

EARLY USER APPRAISAL OF THE MDARS INTERIOR ROBOT

R.T. Laird (SSC-SD), H. R. Everett (SSC-SD), G.A. Gilbreath (SSC-SD),
R.S. Inderieden (SSC-SD), K.J. Grant (CSC)

Space and Naval Warfare Systems Center, San Diego (SSC-SD), CA 92152-7383
Code D371, 53406 Woodward Road, San Diego, CA 92152-7383
laird@spawar.navy.mil, 619-553-3667, <http://www.spawar.navy.mil/robots/>

ABSTRACT

The Mobile Detection Assessment Response System (MDARS) is a joint Army-Navy effort to field interior and exterior autonomous platforms for security and inventory assessment functions at DoD warehouses and storage sites. The MDARS operator interface (host console) is based upon the Multiple Resource Host Architecture (MRHA), a distributed processing system that allows coordinated control of multiple autonomous resources. The current configuration provides control for up to 32 interior and/or exterior robotic vehicles, with provision for near-term integration of remote fixed sensors or sensor pods.

The MDARS-Interior (MDARS-I) effort, initiated in 1989 to improve the effectiveness of a shrinking guard force, was quickly expanded to address the intensive manpower requirements associated with accounting for high-dollar and critical assets. An integral component of the MRHA is the *Product Assessment System*, which extends the capabilities of MDARS into automated inventory management. The robotic platforms are equipped with an RF interrogator for reading interactive transponder tags attached to high-value or critical inventory items. Upon interrogation by the robot, individual tags respond with the stock number and any other relevant information concerning the product. The data collected is stored in a database for comparison with expected conditions, and any discrepancies subsequently flagged by geographic location for investigation.

In 1995, a Broad Agency Announcement (BAA) contract was awarded to Cybermotion, Inc., of Roanoke, VA, to expand upon the government-developed interior prototype by adding improved navigation, extended intrusion detection, and RF interrogation of inventory. The upgraded platform passed formal Technical Feasibility Testing with flying colors in February 1997. Eight interior robots are currently undergoing further evaluation at a number of locations: Building F-36 at SPAWAR Systems Center in San Diego (SSC-SD); the Camp Elliott Warehouse in San Diego; Anniston Army Depot in Alabama; and at the Cybermotion facility in Virginia. This paper provides an overview of the MDARS-I robotic security system, some insights into the problems associated with installation, and a first assessment of lessons learned from the formal Early User Appraisal conducted at Anniston Army Depot.

1. Background

MDARS is a DoD development effort managed by the Product Manager, Physical Security Equipment (PM-PSE), with SSC-SD providing technical direction and MRHA development. The MDARS system, which provides an automated robotic security capability for storage yards, petroleum tank farms, rail yards, and arsenals, includes multiple supervised-autonomous platforms equipped with intrusion

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE APR 1999	2. REPORT TYPE N/A	3. DATES COVERED -			
4. TITLE AND SUBTITLE Early User Appraisal of the Midars Interior Robot		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S) R. T. /Laird; H. R. /Everett; G. A. /Gilbreath; R. S. /Inderieden; K. J. /Grant		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Space and Naval Warfare Systems Center Code D371 53406 Woodward Road San Diego, CA 92152-7383		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 22	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

detection, barrier assessment, and inventory assessment subsystems commanded from an integrated control station.

1.1 Objective and Value of EUA

From a technical perspective, Early User Appraisal (EUA) was intended to test the MDARS-Interior BAA prototype in a real-world warehouse environment under actual operational conditions. The goal was to obtain feedback on system performance from the intended end-user (i.e., the Anniston guard force), so that user-suggested improvements to the MDARS system could be incorporated. Two fundamental guiding principles governed all actions in preparation for EUA:

- The installed system had to satisfy the needs of the user and clearly demonstrate “value added” in a robust and reliable fashion. Simply proving technical feasibility was insufficient.
- The EUA evolution had to be carefully controlled to ensure success. Close coordination with the user was essential, and system complexity had to be optimal and in keeping with the focused objective (i.e., an ultimate 100-percent solution was no more desirable than an insufficient solution).

EUA was an opportunity to determine if the system was ready for operation in a large, semi-structured environment under real-world conditions. EUA also provided an avenue for high-level demonstrations aimed at showcasing MDARS and its component technologies. Finally, data collected and documented during the activities surrounding EUA fed the development of the Request For Proposal for the Engineering and Manufacturing Development Production contract for the MDARS-I system.

1.2 Road Map

Early User Appraisal of the MDARS-I system was performed at Anniston Army Depot (ANAD) in Anniston, Alabama over an eight-month period from March to October, 1998. Site preparation and installation preceding EUA required 13 calendar months from January 1997 to February 1998 (the first site survey to Anniston was conducted in November 1995). EUA technical tasks were carried out by SSC-SD personnel, while logistics and coordination tasks were carried out almost exclusively by Mr. Bob Walker of Computer Sciences Corporation.

Initially, a single robot was installed in the center bay of Building 131 with the temporary Multiple Resource Host Architecture (MRHA) control van located just outside the warehouse). This preliminary setup allowed technicians to troubleshoot the new system efficiently and effectively without adversely impacting personnel working in the guard station. As success in the preliminary stages of site installation was incrementally achieved, the control functions were relocated to Guard Post 7 in Building 53. During the final phase, the MRHA console was integrated into the existing guard station routine and became a working asset for the security staff. A second robot was installed as a backup in case of a failure with the primary unit.

This approach effectively circumvented problems experienced in the past where potential robotic solutions were thrust prematurely into the daily routine of the targeted user, before on-site debugging had been fully completed. Historically, failure to address this issue has created misunderstandings and disenchantment that often could not be reversed, generating in the process bad publicity that a number of

programs were unable to survive. It must be recognized that MDARS is a system, as opposed to a collection of individual components, and a very complex system at that. Lack of appreciation for the full spectrum of demands of effective systems integration has led to the embarrassing failure of more than a few major government programs.



Figure 1. MDARS temporary control van (the big one on left!) adjacent to Building 131 at Anniston Army Depot in Alabama.

2. System Description

2.1 Interior Robotic Platform

The MDARS-I robot is based upon the Cybermotion K2A platform with enhancements resulting from a BAA contract awarded in 1995 (Holland, et. al., 1995). The platform base houses the drive control electronics and processor (DC1) along with the 24-volt battery set. The base uses a three-wheeled concentric-shaft synchro-drive system; all three wheels are driven by the same motor. A separate motor is used to steer the wheels. This design allows the platform to rotate in place with minimal clearance requirements. The platform emergency stop switches are mounted on the outside of the base.

The turret, which mounts directly on top of the base, contains the navigational sensors and recharging receptacle, along with processors for vehicle docking beacon control and security sensor management. Also inside the turret is the high-level MRHA interface processor which runs *Windows NT* and the MDARS-I *Scheduler* software. The *Scheduler* is responsible for controlling the overall operation of the platform and for communications with the MRHA. An ARLAN 640 Ethernet RF modem, adjacent to the *Scheduler PC/104* processor stack, provides the communications link to the MRHA. A narrow-beam sonar skirt attached to the front of the turret along with wide-beam front-, rear-, and side-looking sonar transducers support collision avoidance and navigation (world modeling) functions. The turret is covered by a black plastic hood that contains the Savi RF tag interrogator and the analog audio/video transmitter.

Perched on a vertical boom above the turret is the *Security Patrol Instrumentation (SPI)* subsystem, which consists of an enhanced sensor package used for security and environmental monitoring, and an integral pan/tilt unit that contains an IR-illuminated video camera and a high-gain microphone. An ultraviolet flame detector and the primary intrusion detection sensors (an 8-12-micron passive infrared array and a K-band microwave intrusion sensor) are swept at 360 degrees/second in a rotating dish atop the *SPI*. Directly below the security sensors are a broad-spectrum gas sensor, temperature sensor, humidity sensor, and ambient light sensor, all of which are used for environmental monitoring.

Enhancements made to the *K2A* during the 1995 BAA contract, in addition to technology transferred to Cybermotion under a Cooperative Research and Development Agreement with SSC-SD in years prior to EUA, have resulted in a refined product that Cybermotion now offers commercially, namely the *SR2/ESP*.



Figure 2. MDARS-I robotic platform based upon the Cybermotion K2A.

2.2 Multiple Resource Host Architecture

The MRHA (Figure 3) is a distributed processing system that controls and coordinates the operation of multiple autonomous interior and exterior remote platforms. The system is designed to run automatically with minimal user oversight until an exceptional condition is encountered. This requirement implies the MRHA must be able to respond to exceptional events from several robots simultaneously. Communication among the MRHA processes is based upon Windows Sockets over TCP/IP. This approach permits functionality to be distributed among a number of processors, and allows the physical configuration of the MRHA to change according to the needs of the application/site.

There is no requirement for each process to run on its own dedicated computer as is implied by Figure 3. Multiple processes can run on the same computer given sufficient resources such as CPU horsepower and RAM, although there are few restrictions as to how components can be combined. In addition, components can be executed on remote processors that are connected over a wide-area-network using TCP/IP. The bandwidth requirements for inter-process communications are very low (i.e., < 28.8 Kbps). This distribution of function enables human supervision and interaction at several levels, while the hierarchical design facilitates delegation and assignment of limited human resources to prioritized needs as they arise (Everett, et. al., 1994).

The *Supervisor* process sits at the top of the hierarchy and is responsible for overall system management and coordination. The user interface provides a “big picture” representation of secured areas and system resources. The *Supervisor* has at its disposal a number of process resources, such as one or more *Operator Stations*, two or more *Planner/Dispatchers*, a *Product Assessment Computer*, and a *Link Server* (Laird, et. al., 1993).

User intervention is required only when a platform encounters an exceptional condition, such as an environmental hazard or a security breach. Exceptional conditions are prioritized and an *Operator Station* display is automatically invoked, whereupon the user is informed of the situation and of response options via the *Operator Station* display. This interface allows a user to directly influence the actions of an individual platform, with hands-on control of destination, mode of operation, and camera functions. Also, the display provides detailed operational and diagnostic system information. The *Supervisor* and *Operator Station* displays have been similarly configured to provide the user with consistent user-friendly, graphical interfaces. Both modules support point-and-choose menu buttons for user-selectable options, commands, and navigational waypoints.

The *Planner/Dispatcher* process (an integration of the Cybermotion “Dispatcher” and the SSC-SD “Planner”) is responsible for assembling and downloading paths to platforms.

The *Link Server* provides the interface between the host MRHA and the various robotic or fixed sensor platforms, and maintains a blackboard data structure of platform status information for immediate retrieval by other MRHA resources on the LAN. The IP address of each remote resource is listed in the *Link Server* initialization file. At system startup, the *Link Server* establishes a virtual connection to each remote entity using a UDP/IP Windows Socket on a known port number, or directly using an RS-232 serial port. The physical connection to the remote platform is typically made using a wireless modem attached to the *Link Server*. Most wireless modem networks provide a repeater capability that allows for extended range/coverage in “difficult” environments.

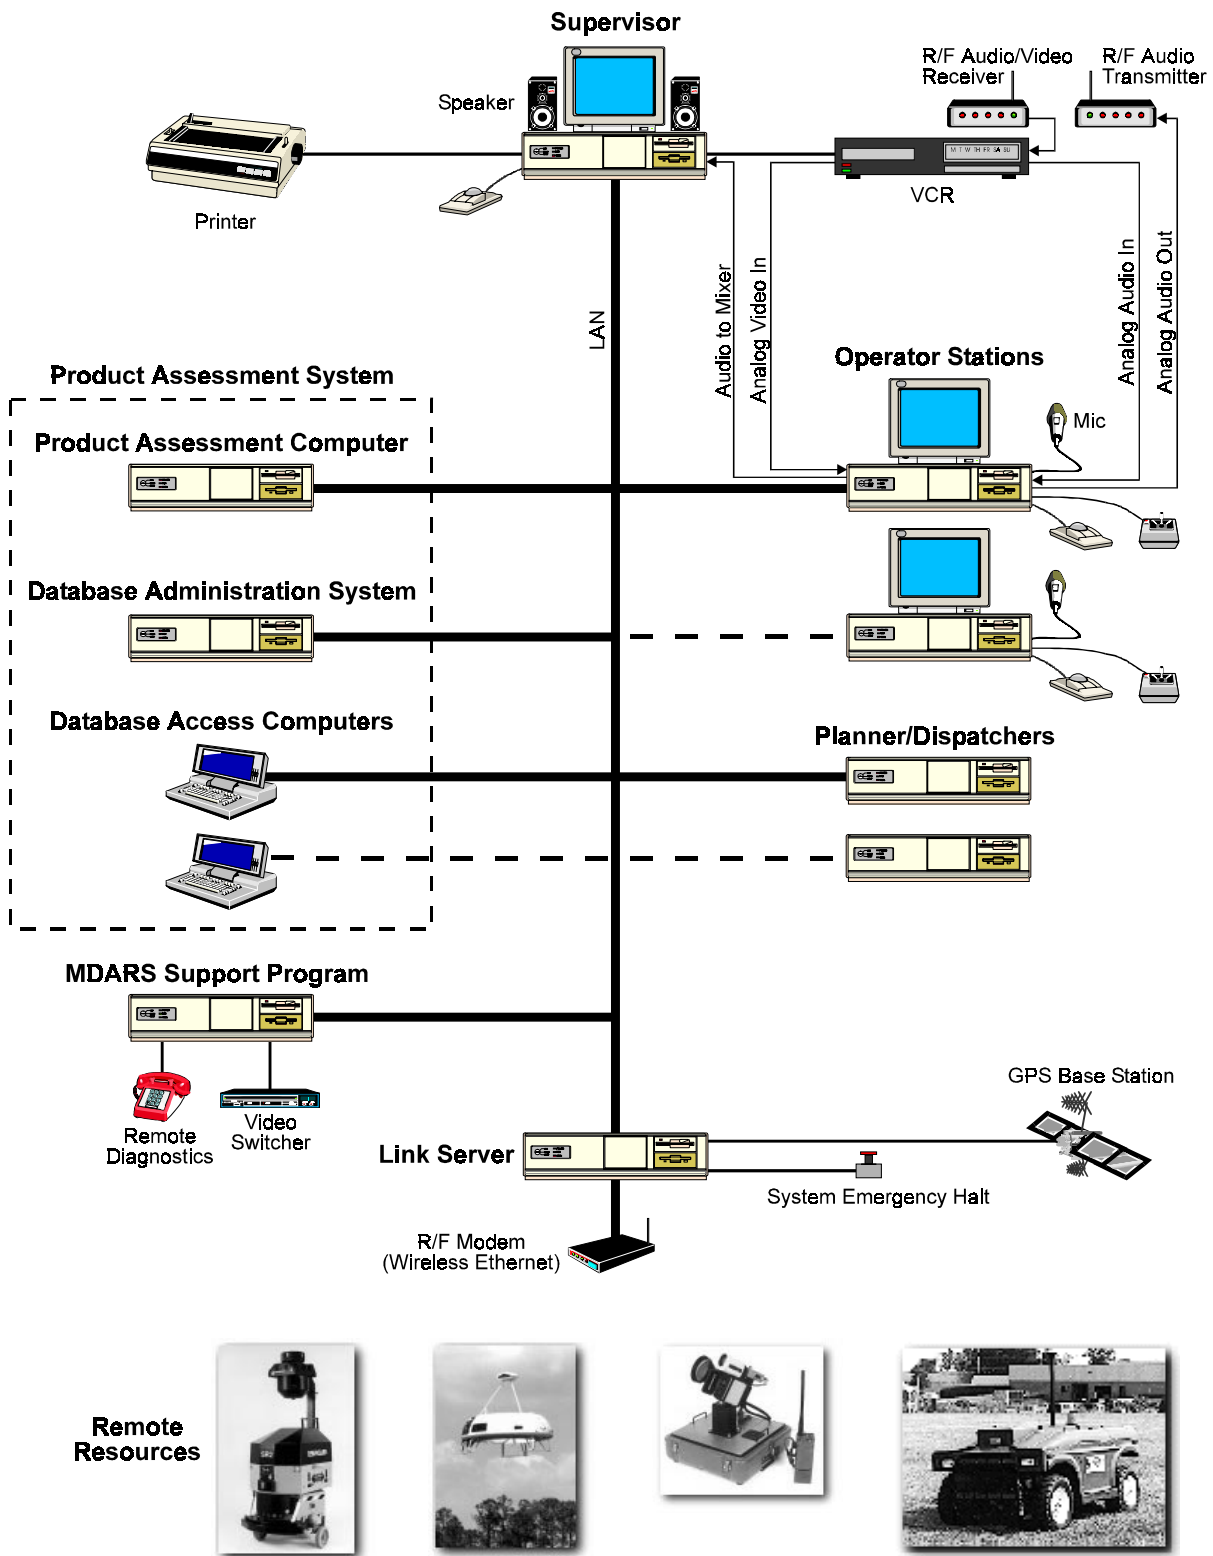


Figure 3. Block diagram of the Multiple Resource Host Architecture (MRHA).

Remote resources are considered slaves to the MRHA, in that they only respond to queries and do not originate communications of any kind. Each remote platform, by definition, is required to process a standard set of functions based upon its logical type or kind. The *Link Server* periodically requests status from the remote platforms, and makes this status information available to other MRHA processes. In addition, the various MRHA processes generate commands (in the form of function-execution requests) that are sent through the *Link Server* to a specific remote platform. Results from an executed command are passed back through the *Link Server* to the originating MRHA process. Commands may be sent from MRHA processes at any rate to the remote platforms and are completely asynchronous. The *Link Server* communicates with the remote platforms in lock-step: it sends a command and then waits for the response (if any) before sending the next command. If a response is not received within a specific time period, then the command is re-transmitted a predefined number of times before marking the remote platform as non-responsive.

2.3 Product Assessment System (PAS)

The PAS maintains a database of high-value inventory, verified by an RF tag reading system onboard the robot, and correlated to geographical location within the warehouse. RFID tags are attached to high-value, controlled, or special-interest items. For EUA, the MDARS-I platforms were equipped with omni-directional Savi Technology *Model CP-1010A-1* RF interrogators. The platforms activate the interrogators at pre-scheduled patrol stops to collect data from all the RFID tags in the vicinity. This information is uploaded to the MRHA, which stores the information in a relational database for item location assessment and reporting.

The *Product Assessment Computer* (PAC) is responsible for uploading RF tag data from each platform, and for transferring the raw data to the *Product Database Computer* (PDC, i.e., database server). The *Database Administration System* (DAS) performs the inventory position-estimation function, or item location assessment, using the raw tag data collected over time. Multiple *Database Access Computers* (DACs) serve as the user-interface to the MDARS database. The DACs provide several database reports and, in the future, can provide an automated interface between the MDARS database and existing site database systems.

Communications between components of the PAS and the database server are based upon the Open Database Connectivity (ODBC) interface using Windows Sockets over TCP/IP. At the application program level, the PAS components make ODBC calls that are translated into structured query language (SQL) and then sent to the database server via a Windows Socket TCP/IP connection. The PAC is the interface between the MRHA and the rest of the PAS, meaning the PAC is the only component of the PAS that connects to the Supervisor as a “recognized” resource on the LAN. The PAC also connects to the database server via a LAN. As this implies, there must be a physical connection between the MRHA LAN and the LAN on which the database server resides, although the two systems (the MRHA and PAS) are typically thought of as separate for political purposes. The DAS and the DACs connect directly (and only) to the database server, and have no knowledge of the rest of the MRHA.

The PAC, DAS, and DACs interface with each other via database tables managed by the database server. The PAC creates inventory survey tables that the DAS periodically processes. The DACs display data that has been processed by the DAS and allow for manual data entry by inventory personnel.

Item location assessment is performed by the DAS to determine the perceived location of tagged items. Knowing where an item *is*, as opposed to where it's *expected* to be, is of great value to inventory personnel. The DAS determines the perceived location of each item by triangulation, using the three strongest readings from a series of robot-performed interrogations. The Savi interrogator provides signal strength (which is inversely proportional to the square of the distance from the interrogator) for each tag reading. The calculation of perceived location uses this signal strength as a weighting factor in its location estimate: the perceived location is computed to be closer to those positions from which higher signal strengths were received than to those positions where signal strengths were lower. Improving the ability to accurately assess the location of each tagged item is an area of ongoing development for SSC-SD engineers.

The DACs allow inventory personnel to access item data and get regular or *ad hoc* reports on items that are being tracked. Typical output is in the form of exception reports highlighting potential problems, such as items perceived to be missing from the facility or moved from their expected locations. Personnel can also use DACs to locate specific items using perceived and expected locations. Report and specific-item search information is available in both text and map form.

3. Installation

3.1 Logistics and Coordination

The importance of employing a competent logistician/expediter during the installation process cannot be overemphasized. Such an individual is invaluable in terms of coordinating the activities surrounding system installation. The closer (physically) the logistics team is to the gaining site the better, as frequent site visits are required to meet with local facilities and base personnel to coordinate tasking. The MDARS logistician was responsible for coordinating nearly every activity that involved more than one agency. This was a monumental task, considering the nature of the mission, which was inherently intrusive and potentially very disruptive. Hundreds of man-hours were saved by assigning (i.e., sacrificing) a single individual to the administrative/logistical/political needs of the program.

Points of contact for the security and inventory (warehouse) functions were established well before installation began, and informed of all tasks that might impact business within the warehouse or guard station. In addition, a memorandum of agreement was signed by PM-PSE and ANAD that outlined the responsibilities of all parties involved in the system installation and operation.

3.2 Site Preparation

Site preparation began in January 1997 after several months of waiting for final approval of ANAD as the EUA site. A number of field trips were made by SSC-SD personnel to install support equipment for MDARS. The actual time spent performing site preparation tasks was approximately five weeks, executed over a six-month period. Progress was slightly hampered by the non-availability of materials, as Anniston is a relatively small and isolated community. Most building materials required for site preparation were available locally, but the majority of computer-related equipment had to be procured in San Diego and shipped to ANAD, requiring additional time that had not been scheduled. In addition, as

ANAD is an operational depot, limited access to facilities and utilities required preparations to be performed during designated hours only.

The goal of the site preparation phase was to install the basic utilities required at the control van, warehouse, and guard post before system installation began. The major site preparation tasks included: running AC power to devices within the warehouse rafters; pulling fiberoptic cable between the guard post, control van, and warehouse; and installing commercial phone lines into the control van and warehouse. As much of this work as possible was performed by local public works personnel to minimize the impact on the engineering team.

AC power within the World War II-era warehouse tended to be very “dirty.” Service was run from the warehouse main breaker panel to a UPS housed within an environmental enclosure mounted to a wall in the MDARS robot recharging area. Power to all other MDARS devices within the warehouse (except the robot chargers) was taken from this UPS. The UPS provided “clean” power and offered significant protection from power surges experienced during lightning strikes, which were common in that area. The environmental enclosure (referred to as the Hoffman box) also housed a *10Base-T* Ethernet hub, a multi-channel video switcher, an audio/video distribution amplifier, and a fiberoptic multimedia extender. Figure 4 shows the MDARS robot charging area within ANAD Building 131 with the Hoffman box in the background.

Shortly before ANAD was chosen as the EUA site, fiberoptic cable had been run between most of the warehouses on the depot, with several spare fibers in each bundle. Consequently, additional fiber for MDARS had to be pulled only a short distance between an adjacent warehouse and Building 131, which saved considerable time and money. Once fiber was run to Building 131, connectivity to Building 53 (the guard post) was achieved.




Figure 4. Hoffman box mounted on brick wall above desk on left in the MDARS robot recharging area during the site preparation phase of EUA.

Four commercial telephone lines were installed within Building 131 for remote connectivity to SSC-SD, yet the number was occasionally insufficient to support the needs of the MDARS development team. One line was used to transmit remote video back to San Diego; one line was dedicated to the MDARS

Support Program (MSP) for remote diagnostics; one line was used as a remote Ethernet bridge between San Diego and Anniston; and one line was used for voice. During system integration at ANAD, the voice line was frequently used to access the Internet for software driver and file downloads. A direct connection to the Internet would have facilitated system development and minimized the need for multiple phone lines. (Toward the end of EUA, a high-speed Internet connection was identified within Building 131 and used for large file transfers.)

Site preparation culminated with the receipt and installation of the temporary control van that was situated adjacent to Building 131. Fiberoptic cable (two pair), the four phone lines, and AC power were run from the warehouse to the control van, which was networked via fiber and a telephone Ethernet bridge to Building 131, the guard post, and SSC-SD.

3.3 System Installation

Site preparation provided the basic utilities needed to begin formal system installation in June 1997, which in turn required well in excess of 500 man-hours over a period of approximately nine months. The major activities included: 1) installation of the temporary MDARS control console within the control van adjacent to Building 131; 2) installation of the robotic platforms and support equipment within the warehouse; and 3) installation of the primary MDARS control console within Building 53 (the Guard Post).

The temporary control van (a rented trailer) was uninhabitable as delivered, and had to be completely refurbished with new carpet, sub-floor cabling, a wireless alarm system, and additional circuit breakers to support the MDARS control console equipment. Figure 5 shows the control van interior during the early stages of system installation.



Figure 5. MDARS control van interior during System Installation Phase of EUA.

A two-bay modular console (Figure 6) was installed near the center of the van for use as the MDARS control station during the preliminary stages of operation and during EUA demonstrations. The console housed the *Supervisor* and *Operator* computer monitors, in addition to the emergency halt switch and the stereo audio speakers. The *Supervisor* and *Operator* mice, along with the *Operator* joystick, sat atop the console desktop. Fixed-camera warehouse video and robot video were displayed on two separate monitors positioned on top of the console. The control van also housed the 19-inch equipment rack that contained the MRHA computers and peripherals. The PAS database server was also located in the control van.



Figure 6. MDARS control station console housing the Supervisor and Operator Stations.

Within the warehouse, a number of tasks were performed during system installation in preparation for actual operation. Nine surveillance cameras were installed in the rafters of the center bay of Building 131, and their outputs were run to a video switcher located in the Hoffman box. The cameras covered each of the aisles within the center bay of the warehouse so that the movements of the robots could be monitored from within the MDARS control van. An RF audio/video receiver was installed near the north end of the center bay of Building 131, with its output connected to an audio/video amplifier located in the Hoffman box. This receiver provided the audio/video uplink from the robot to the control console. Three ARLAN 640 RF Ethernet modems were installed along the center aisle within the rafters of the center bay of Building 131, and connected to a *10Base-T* Ethernet hub located in the Hoffman box. An intercom system was installed for communications between the control van, the MDARS recharging area, and the PAS launch station. Two robot-charging stations, each equipped with an optical docking beacon for platform re-referencing, were placed against the west wall of the MDARS recharging area.

The MDARS PAS launch station (for programming the RFID tags) was set up in the first bay of Building 131. The launch station consisted of a desktop computer system configured as a *Database*

Access Computer, along with a barcode reader, attenuated tag interrogator, and laser printer. AC power, Ethernet, and intercom cables were run to the launch station, with a separate UPS to provide clean power.

After a majority of the support equipment was operational, the robot *virtual-path* programs were developed. This task, referred to as “path installation,” consists of writing a sequence of instructions that direct the robot’s movement within the warehouse using navigational landmarks such as walls or crates for reference. As the process is typically very tedious and time consuming, a minimal-path network was initially installed to validate system operation as quickly as possible. Later, after the warehouse configuration had stabilized, the final paths were added.

In parallel with work being performed in Building 131, a second console was installed in Guard Post 7 located within Building 53. (Actual operation would eventually transition from the temporary control van near Building 131 to Guard Post 7.) The guard post installation differed significantly from the MDARS control van configuration in that all of the computer and peripheral equipment in the former installation was housed within the console itself (Figure 7). The 19-inch equipment rack was eliminated and replaced by a small debug station consisting of a computer monitor and a keyboard located on a mobile stand. In preparation for the console installation, dedicated AC circuits were run from the main breaker panel in Building 53 to Guard Post 7. Fiberoptic cable was pulled from the telephone closet in Building 53 (where it terminated its run from Building 131) to Guard Post 7. Four additional telephone lines were installed near the debug station, and dedicated to the same functions as their counterparts in the control van.



Figure 7. Final MDARS guard post console (white cabinet on right) in Building 53.

Installation of the guard post control console was significantly easier than the control van console installation as nearly all of the equipment was co-located, making cable runs much simpler. The MDARS team was also becoming much more efficient at configuring the computer and support equipment, but preparation and installation still required in excess of 75 man-hours. Limited (after-hours-only) access to the guard post hindered progress, especially considering that the team had already put in a full day at the warehouse before shifting to the guard post. The confines of the guard post made

for restricted working conditions such that a small contingency was all that could be assigned to the task. The guard post console became officially operational in March of 1997, which marked the formal beginning of EUA.

4. Operation

The system was actually up and running several months prior to the formal kick-off, as the MDARS team had to first validate operation before handing control over to the guard force. Several software bugs, mostly related to reliable navigation, had to be corrected prior to transfer to the user. Naturally, the amount of time required to correct these perceived system deficiencies was considerable. Reliability was paramount, yet extremely difficult to achieve.

Several months were spent performing final system integration and testing in preparation for the pending EUA demonstrations and the formal transition to the guard force. Limited access to the warehouse during relatively short trips (in terms of the number of days spent on travel) resulted in rushed efforts to solve non-trivial technical problems under considerable pressure (see Section 5). Consequently, system installation and final integration took longer than anticipated, which minimized the amount of time available for debugging before formal EUA activities began. Once the MRHA was completely operational, the path programs had been fine-tuned, and all of the major hardware problems had been solved, the system ran quite smoothly.

4.1 *Controlling Multiple Robots from a Single Console*

The system was operated by engineers from the temporary control van adjacent to the warehouse. A single robot patrolled the center bay during duty hours only, and was parked on the charger after-hours when the warehouse was secured with the fixed IDS (motion detection) sensors activated. A second robot was later installed to provide complete coverage over weekends when continuous patrols were required. This second robot also served as a backup in case of a failure on the primary platform, and was used during VIP demonstrations to perform continuous inventory assessment runs.

All aspects of system operation could be monitored from the control van, with both fixed-camera video and robot video allowing the user to see exactly where the robot was, in addition to what the robot was seeing. The *MDARS Support Program* (MSP) would automatically switch the fixed-camera video in order to track the movement of a selected platform as it patrolled. Audio from the robot was output at the console so that the user could hear within the warehouse. Direct communications with MDARS personnel within the warehouse were provided via a hard-wired intercom system. The MRHA displayed the state and position of the robots as they patrolled, and all of the functions were at the disposal of the user to perform the dual missions of physical security and inventory assessment.

From a remote site such as SSC-SD in San Diego, developers could monitor the status of the system at ANAD using video from a *GYJR* remote video transmitter and the MSP. Fixed-camera video and robot video were both connected to the *GYJR* unit so that the remote site could switch between either video source. The remote site could also call into the MSP at ANAD by telephone to obtain system status, and manually switch between the fixed-video cameras in order to visually assess a situation. Additionally, using the telephone Ethernet bridge established by the 3COM Access Builder, the remote site could take

complete control of the system at ANAD and operate it remotely. Such capabilities saved considerable money by allowing the MDARS team to support and maintain the system remotely from San Diego.

The Anniston system was typically configured with the primary robot performing directed patrols and security sweeps under *Operator Station* control, while the secondary robot performed continuous inventory assessment runs autonomously under *Supervisor* control. Multiple robot simulators were also run to test operation with up to four platforms, the scenario for all of the EUA demonstrations.



Figure 8. Seven major demonstrations involving two real robots and two simulated robots were conducted during EUA for visiting officials and potential users of the MDARS Interior system.

4.2 User Training in Preparation for Turnover

Immediately prior to system turnover to the local guard force, both supervisory- and guard-level training were given to ANAD personnel by SSC-SD. The supervisory-level training included system configuration and maintenance, as well as operational training. This was a three-hour training session that should have been done over a period of perhaps two weeks with extensive hands-on exercises. The guard-level training covered only the operational aspects of the system and was performed in about two hours. Figure 9 shows one of the guard-level training sessions performed at the MDARS control console in the temporary control van. The guards were given the opportunity to perform hands-on exercises which were critical to their understanding of the system's capabilities. As the guards were the primary users, they should have received considerably more training. Insufficient resources were allocated to training at all levels, however, and consequently the users were not adequately prepared to use the system as effectively as they could have to enhance their security or inventory needs.



Figure 9. Guard-level MRHA training performed at the remodeled MDARS control van near Building 131.

Based upon feedback from the guard-level training, a number of changes were made to the MRHA and man-machine interface. For example, the site map that is displayed on the *Supervisor* and *Operator Station* monitors originally consisted of a plan-view drawing of Building 131. The guards informed SSC-SD personnel that they did not know the exact location of the building. They requested that we include the surrounding buildings in the map for reference, and that we add a true north indicator so that they could orient themselves properly when viewing the MRHA map displays.

4.3 PAS Tests

In order to assess the feasibility of using RFID for inventory tracking in DoD storage facilities, some fairly extensive evaluations were performed during EUA using the Savi *TyTag*. Several types of tests were conducted in an attempt to answer a number of technical questions.

Question: “How well would the tags be seen in this environment of densely packed crates containing a heavy concentration of metal objects?”

Answer: These tests consisted of attaching tags to the exterior surface of wooden crates filled with M-16 rifles, in progressively more hidden positions relative to the interrogator. In the worst-case scenario, tags were inserted between crates that were stacked three high and six or seven deep, with the furthest tag at least 26 feet from the interrogator. The results of these tests showed very good visibility of the tags: in the worst-case scenario, 50 of 52 (or 96 percent of) tags were seen, with the two missed tags being furthest from the interrogator (i.e., 26 feet away).

Question: “How well would the perceived locations of tags be assessed?”

Answer: 117 tags were distributed in the warehouse, and the robot performed location estimation on all tags using normal patrol routes. Results showed an average difference of about 15 feet between actual and assessed location, well within the MDARS target specification.

Question: “Would there be any problem with reading densely packed tags?”

Answer: A large number of tags (73) were packed in a box, about one inch apart, all oriented in the same direction. The robot-mounted interrogator read 100 percent of the tags from about six feet away.

Question: “Would tags need to be placed at certain orientations relative to the interrogator to be read properly?”

Answer: Tags were situated in the patrol area, first parallel to the robot path then perpendicular to the path. The results showed slight improvement (10 percent) in tag readability when tags were arranged perpendicular to the path versus parallel, but overall results were very similar.

The MDARS *Product Assessment System* can displace a large number of expensive fixed interrogators in that a single robot carrying one interrogator can patrol an area on the order of 20 football fields in size. The mobility and maneuverability of an interrogator allow for flexibility in interrogation patterns, and can also potentially facilitate the detection of shorter-range (less-expensive) RFID tags.

5. Problems and Solutions

The following are a few representational examples of the types of problems (and their fixes) encountered in the real-world operational warehouse environment at ANAD.

5.1 Navigation (Robot Lost)

The robot depends on the presence of walls (or wall-like structures such as crates, racks, boxes, etc.) or retroreflective tape markers to navigate in the warehouse. Each path downloaded to the robot contains information about the location of these navigational cues so the robot can use its sensors at the appropriate times to detect these features and correct its X-Y location and heading. If cues are missed, either due to sensor limitations or changes in the operating environment, or if an incorrect feature is misinterpreted as a cue, the robot’s perceived location will begin to differ from the actual location. This can eventually result in the platform becoming blocked by what it perceives as an obstacle.

In one situation at Anniston, some crates were inadvertently moved about six inches to a foot further away from the aisle during normal warehouse operations (Figure 10). The robot still saw the crates, but they were further away than expected, which caused the robot to angle in toward the crates to correct its offset. The crates that followed the displaced crates were in their normal position, but were misinterpreted as obstacles because the robot’s position was now incorrect. The solution in this case was to have the warehouse personnel move the offending crates back into line with the rest of the crates.



Figure 10. Sloppy arrangement of crates along an aisle (the white line on the ground marks the aisle boundary).

In another situation, a retro-reflective stripe was placed on a warehouse support column to serve as a navigational cue. Normally this would be a fairly robust cue since the post theoretically will not move. In this case, however, a forklift backed into the post, rudely displacing it by about five inches (Figure 11). This caused the platform to experience navigational errors, since the post no longer appeared to be at its expected location. The initial solution to this problem was to reposition the stripe on the post until the post was restored to its original configuration (at the gentle urging of several large men wielding sledgehammers).



Figure 11. Retro-reflective tape marker (top, center) on displaced support column after being struck by fork truck.

5.2 Cracks

Concrete expansion joints, as well as cracks caused by settling and thermal expansion, are a source of two different navigational problems. The first of these occurs due to the high acceleration forces placed on the entire platform when a lateral crack is traversed. Since the vehicle lacks any kind of shock absorption mechanism (except for the SPI boom), any floor irregularity encountered by the wheels is immediately coupled to the rest of the platform. When a crack is traversed, the jarring experienced by the platform is significant. This violent movement causes the safety bumper to vibrate up and down, which occasionally causes the tactile switches inside the bumper arms to make contact, stopping the platform. This in turn causes an *emergency-halt error* to be reported by the MRHA, whereupon the platform is assigned to the *Operator Station*. The user must then hit the *Resume* button and release the platform to continue normal patrols. This problem was partially solved by decreasing the sensitivity of the bumper switches and filling in some of the more serious cracks.

The second problem occurs when one of the wheels falls into a longitudinal crack, usually an expansion joint, running parallel to the path. When the platform attempts to turn, either for a navigational correction or to change paths, the trapped wheel is inhibited from rotating because it is partially embedded in the crack. This causes the DC-1 drive controller to increase the current to the steering motor, often resulting in an over-current condition, whereupon the platform halts with a *current-limit* status code. This situation occurs more frequently as the battery discharges, since more current is required to achieve the same amount of power with less voltage.

Once the vehicle has been halted in this fashion, simply resuming the path will sometimes get it out of the crack. Other times the vehicle can be manually steered out of the crack by using telereflexive mode. However, in some situations, particularly when the battery is low, the vehicle can only be recovered using the drive pendant and/or by physically lifting or pushing it out of the crack. Another serious side effect of this longitudinal crack problem is that a wheel will get wrenched out of alignment by the resulting torque, with significantly degraded dead-reckoning capability as a consequence.

There are two potential solutions to the longitudinal crack problem: 1) adjust the paths to keep them away from the cracks, and 2) fill the cracks with concrete patch. A newer version of the Cybermotion robot is available (the K3A) that uses a dual-wheel design that may eliminate or at least alleviate this problem significantly.

5.3 SPI False Alarms

The SPI uses a rotating sensor array to scan the surrounding area for intruders. An essential part of this system is a slip ring, which is basically a set of fixed brushes that make contact with circular rings on the rotating element of the SPI. This allows the spinning sensors to rotate freely without being hindered by connecting wires. After operating reliably for several months at Anniston, the SPI began to report intruders with a high degree of certainty when no intruders were present. This problem was eventually tracked down to accumulated brush dust. As the brushes slide along the metal rings of the rotor, a small amount of brush dust is generated from normal wear, and over time enough falls onto the PC board below to cause intermittent shorting between the brush mounts. The solution was to coat the solder joint as well as the non-contact portion of the brushes with a non-conductive conformal coating.

5.4 Robots Stuck on the Charger

There were three situations in which the robot could get “stuck” on the charger, and in all three cases the robot had been automatically sent to the charger due to a low-battery condition. When the robot docks with the charger and sees that it must recharge, it continues to run the most recent path program, which monitors the charge state of the batteries. Once the batteries are fully charged, the program resets the low-battery status and terminates. The MRHA sees that the platform is idle and that the low-battery status is no longer present, and returns the robot to random patrols.

In two of the problem cases, the recharging robot had been assigned to the *Operator Station* by a guard and then halted via the *Halt* button (no one knows why the user halted the robot while it was docked to the charger). This unnecessary action caused the recharging program to terminate prematurely, leaving the robot in a bad state with the low-battery status still set, but not running a program, and with a path-interrupted status. If left on the charger in this state, the batteries will eventually recharge, but since the recharging program is not running, the low-battery status will never be reset.

There are two possible near-term solutions. First, pressing the *Resume* button on the *Operator Station* will restart the charge monitoring program, which will terminate once it detects that the batteries are fully charged. The second solution is to re-reference the robot using the *Reference* button on the *Operator*. However, since the low-battery status will still be set, the *Supervisor* will immediately send the robot back to the charger as soon as it is released by the *Operator*. Once the batteries are recharged, however, the robot will return to random patrols. The long-term fix is to disable the *Halt* button while the robot is charging at its dock.

The third situation occurred due to a contact failure with the charger, probably because the robot was not fully mated to the recharging prong. There are two sensors used for docking and recharging: one is the “contact” sensor that indicates to the robot when it has made contact with the recharging prong, while the other is the “float” sensor that indicates when the batteries are fully charged. During the initial docking maneuver, only the contact sensor is monitored, but once contact has been made, only the float sensor is monitored.

Unfortunately, when contact with the charging prong is lost, the float sensor immediately goes high, indicating (rightly or wrongly) that the batteries are charged. The recharging program consequently terminated and the robot was placed back on random patrols. Since the robot had just docked with the charger, the low-battery status quickly returned, and the *Supervisor* sent the robot right back to the charger. Although not yet addressed, one possible solution is to monitor the *contact* as well as the *float* sensor. If *float* goes high but *contact* is lost, then the robot should jog forward slightly to regain contact with the charger.

5.5 Script Scheduling

Script files are sequences of commands for a specific platform to perform various duties in a sequential fashion. Examples of these commands include putting the platform in *Inventory* mode, and then sending it to various locations within the facility. Scripts can be executed “on the fly” or scheduled via the initialization file to occur either at system startup, at a specific time, or some elapsed time period after startup. If there is no script scheduled for a platform, then the platform is available for random patrols, unless the platform is offline. The original intent for scripts was to be able to specify a particular patrol

route for the platform, either for intruder detection or inventory assessment. More advanced emergent usage of the script capability during EUA exposed weaknesses in the script execution implementation.

The first problem was situational. The Anniston facility used two platforms to patrol a single warehouse, so one platform could patrol while the other one remained at the dock recharging its batteries. At startup, the system would not know which platform was patrolling and which platform was waiting to patrol. The schedule would be preset in the initialization file to put a specific platform on patrol, and place the other platform off-line at the charger. After a few hours, the patrolling platform would be instructed to go to its charger and go off-line, and the resting platform would go on patrol. (The docking paths are set up such that only one platform at a time can be running in that area of the warehouse.) The problem occurred when the first platform to patrol was not previously fully charged, and would run low on battery power during its patrol time. The system detected the low-battery condition and dutifully sent the platform to the dock to be recharged. Unfortunately, there was no method for releasing the second platform as a backup patrol (before its pre-scheduled release time). So the patrol area would be left unsecured with both platforms at the chargers.

The second problem was synchronization. A single script file could only control a single platform, and there was no way for a script to know what another platform was doing at a given time. For example, when the system needed to put the patrolling platform back on the charger and place the second platform on patrol, the steps performed had to occur in a given sequence to be successful. The patrolling platform must have reached its dock before the waiting platform could proceed to the patrol area. The only way to accomplish this was to schedule a significant time gap between the two scripts (i.e., one sending the patrolling platform to its dock, and the other placing the second platform on patrol). For instance, if the patrolling platform was at the farthest point from the dock when it got the command to return to the charger, it could take several minutes for the platform to traverse the entire path. In addition, when a script was scheduled to execute, the patrolling platform may have been on a path which could take a long time to complete, including security-survey or inventory-read stops, increasing the likelihood of a scheduling conflict.

These problems were addressed by redesigning the scripting system to add more sophisticated capabilities. Now, a single script file can issue commands to multiple platforms, and the script will know when each command has been executed so it can synchronize activities between platforms. For instance, the script may contain the commands “DOCK 1” followed by “RANDOM 2:00 2,” indicating platform 1 should proceed to its charging location, and then platform 2 should perform random patrols for two hours. The RANDOM instruction will not be executed until the DOCK instruction has been completed. A limited set of control flow commands were introduced, including a basic loop, conditional testing and branching, as well as CALLing or CHAINing other script files. The script file can now specifically direct and coordinate platform activities within the same patrol area.

In the new design, each patrol area will have its own script file to direct and coordinate activities within that specific zone. This approach allows users to control one patrol area very differently from another, but easily coordinate activities within a common zone. The script file will have “exception handlers” to indicate what actions should be taken when specific situations are detected. For instance, if two platforms are to alternately patrol an area, and the patrolling platform has a low battery, the script would contain instructions to send the patrolling platform to the dock and then place the other platform on patrol. The script processor will also allow the use of variables to simplify scripts and make them more generic.

6. Conclusion

The Anniston MDARS-I system remained up and running for five months after the last VIP demonstration was conducted in early June of 1998. Relatively speaking, there were minimal problems. Some security guards ran the system more than others (some refused to run it at all), but quite a bit was learned from the comments and feedback provided in either case. All the design deficiencies uncovered in the process were formally addressed with Engineering Change Proposals and Trouble Reports. SSC-SD monitored the system remotely from San Diego, with an e-mail tie-in to supervisory personnel within the security organization at ANAD. From a technical perspective, a few specific observations are worth noting in retrospect.

First and foremost, the MRHA software did not crash once during the entire EUA time frame, which given its complexity and the accelerated pace of development (including both a language conversion and incorporation of a new operating system), was quite an impressive feat in and of itself. Secondly, Cybermotion's SPI security sensor subsystem, required to detect intruders over a 360-degree field-of-view out to 10 meters, consistently demonstrated reliable detection to 30 meters, and in some cases out to 100-130 meters. Thirdly, the Product Assessment System likewise surpassed all expectations, finding RFID tags buried deep within rows of stacked crates filled with high-metal-content inventory (i.e., rifles), and resolving their position to within the required specification of 15 feet.

And last but certainly not least, thanks to years of characterizing and fine-tuning (at the beta-test warehouse at Camp Elliott in San Diego) with close cooperation from Cybermotion, reliable autonomous navigation in an unstructured warehouse environment was clearly achieved. (The BAA requirement called for not more than one navigational failure per platform per shift.) There were only three situations following turnover where a robot did get lost, and two cases were due to excessive clutter left in the path by warehouse personnel. The third scenario resulted from a prolonged loss of electrical power during a storm, resulting in the robot's battery going dead before the MRHA recovered. (Production systems will be equipped with an eight-hour emergency backup via the UPS, and have a more stringent requirement for not more than one "platform-lost" incident per month).

In closing, the post-turnover trials of EUA proved to be an invaluable evolution for the MDARS development team, at essentially no additional cost to the program, but unfortunately had to be terminated in November 1998 for programmatic reasons. The lessons learned and feedback received weighed heavily in the structuring of the Request For Proposals (RFP) for the follow-on Engineering Model Development (EMD) contract.

7. References

- Everett, H.R., Gilbreath, G.A., Heath-Pastore, T.A., Laird, R.T., "Controlling Multiple Security Robots in a Warehouse Environment," AIAA/NASA Conference on Intelligent Robots, Houston, TX, March 1994.
- Everett, H.R., Laird, R.T., Thomas, K.M., Boroughs, R., "Site Installation Plan/Report for the Mobile Detection Assessment Response System (MDARS) Early User Appraisal (EUA)," SSC-SD, TD 3038, San Diego, CA, August, 1998.
- Holland, J.M., Martin, A., Smurlo, R.P., Everett, H.R., "MDARS Interior Platform," Association of Unmanned Vehicle Systems, 22nd Annual Technical Symposium and Exhibition (AUVS '95), Washington, DC, July, 1995.

Laird, R.T., Everett, H.R., Gilbreath, G.A., "A Host Architecture for Multiple Robot Control," ANS Fifth Topical Meeting on Robotics and Remote Handling, Knoxville, TN, April, 1993.