A Multisensor, Vehicle-Mounted, Teleoperated Mine Detector With Data Fusion

John McFee[a], Vic Aitken[b], Robert Chesney[a], Yogadhish Das[a], Kevin Russell[a]

[a]Defence Research Establishment Suffield,
Box 4000, Medicine Hat, AB, Canada, T1A 8K6

[b]presently with Thomson Technology Ltd.,
250 Fielding Road, Lively, ON, Canada, P3Y 1L6

ABSTRACT

The Improved Landmine Detector Project (ILDP) was initiated in Autumn 1994 to develop a prototype teleoperated vehicle mounted mine detector for low metal content and nonmetallic mines to meet the Canadian requirements for rear area mine clearance in combat situations and peacekeeping on roads and tracks. The relatively relaxed requirements, such as low speed and reduced detectability of completely nonmetallic mines, greatly increase the likelihood of success. The ILDP system consists of a unique teleoperated vehicle carrying a forward looking infrared imager, a 3 m wide down-looking highly sensitive electromagnetic induction detector and a 3 m wide down-looking ground probing radar, which all scan the ground in front of the vehicle. Scanning sensor information is combined using a suite of navigation sensors and custom designed navigation, spatial correspondence and data fusion algorithms. Suspect targets are then confirmed by a thermal neutron analysis detector. A key element to the success of the system is the combination of sensor information. This requires coordinated communication between the sensors and navigation system and well designed sensor co-registration, spatial correspondence and data fusion methodologies. These complex tasks are discussed in detail. The advanced development model was completed in October 1997 and testing and improvements are ongoing. Results of system performance during extensive field trials are presented. A follow-on project has been initiated to build four to six production units for the Canadian Forces by the year 2000.

Keywords: Mine Detection, Multisensor, Vehicle-mounted, Data Fusion

1. INTRODUCTION

A comprehensive 1994 report[1] by the Canadian Defence Research and Development Branch (DRDB), based on twenty years of mine detection research and development and analyses of the latest developments in related technologies, concluded that a single mine detection technology that could detect all types of mines under all conditions would be very expensive to develop and would be technically very risky. This is evidenced by the fact that after nearly fifty years and an international investment of hundreds of millions of dollars in R&D, there is still no fielded effective detector for nonmetallic mines. Further, it recommended that the Canadian Department of National Defence should focus R&D efforts on a niche that was of relevance to the CF, was potentially solvable, used the broad technical spectrum available within DRDB and was within the DRDB's budget. Such niches exist for Canada, specifically some mine detection scenarios in rear area combat situations and Operations Other Than War (OOTW). OOTW refers to peacekeeping, demining or operations which do not involve high intensity conflict. Mine detection for these roles is a potentially solvable problem because the requirements imposed on a detector for this role are substantially relaxed compared to those for the general high intensity conflict role. Requirements for the latter include that the detector must operate at high speed without stopping, must have the ability to operate on all terrain, in all weather, at any time of day or in any season, must detect all nonmetallic mines and must detect mines of all sizes. These impose severe constraints on candidate detectors which lead to high expense, technical complexity and risk of failure. A lower risk and cost approach can be taken in rear area clearance or OOTW detection by relaxing these requirements.

Other author information: (Send correspondence to J.M.)
J.M.,R.C.,Y.D.,D.H.,K.R.: Email: John.McFee@dres.dnd.ca; Telephone: 403-544-4739; Fax: 403-544-4704
V.A.: Telephone: 705-682-1517; Fax: 705-682-3358
Supported by the Defence Research and Development Branch, National Defence, Canada, Projects 2HB and 2HE.

The detector can operate at low speed and stopping is allowed. In many situations, the mission has the flexibility to restrict the detector to roads and tracks at limited times of day and in limited environmental conditions. Reduced detectability of completely nonmetallic mines is allowed. This is not a serious restriction since $> 99\%$ of all mines are estimated to have detectable amounts of metal. The detector is required to detect only antitank (AT) and large antipersonnel (AP) mines.

As a result of that study, the Improved Landmine Detector Project (ILDP) was started in Autumn 1994 to design and build an advanced development model of a teleoperated vehicle mounted mine detector for low metal content and nonmetallic mines to meet the Canadian requirements for peacekeeping on roads and tracks. The Defence Research Establishment Suffield was chosen to be the lead laboratory. The approach taken was to employ multiple detectors based on technologies which had previously failed for the high intensity conflict problem or in a single sensor role, chiefly because of high false alarm rates. The output of these detectors would then be combined using data fusion to reduce individual detector false alarm rates and provide redundancy. A teleoperated platform was chosen to improve safety to the operators and the platform was custom designed to have a low signature, in particular ground pressure, with respect to AT mine fuzes to increase system survivability.

2. THE ILDP LANDMINE DETECTION SYSTEM

2.1. Selection of Detectors

Selection of detectors for the system was based on a long list of criteria which had been established in Ref. 1. In fact, the initial selection had already been made in that report. The most significant criteria were a high probability of detection, relatively low technical risk, relatively low cost and suitability for data fusion. A main condition for the latter is that the sensors must ideally have independent outputs, so that multiple alarms due to a target would be indicative of a mine, even though any single detector could occasionally alarm for a non-mine. The five detectors chosen were electromagnetic induction (EMI), infrared imaging (IR), visible wavelength imaging (VIS), ground penetrating radar (GPR) and thermal neutron analysis (TNA). The theory of operation of these sensors is described in Ref. 2.

2.2. System Description

A system block diagram is shown in Fig. 1. The system consists of three major subsystems; the remote detection vehicle which contains the scanning detectors, an attached confirmation detector trailer which contains the TNA detector and a vehicle control station which is housed in a separate vehicle.

A photograph of the system is shown in Fig. 2. The remote detector vehicle (RDV) is an eight wheeled, hydrostatically powered vehicle which was custom-built to DRDB specifications by Q Sine of Calgary. Each pair of adjacent wheels is mounted on a rocker arm to minimize rolling and pitching of the vehicle and to maintain a constant ground pressure as the wheels roll over obstacles. The vehicle provides electrical and hydraulic power to all other subsystems. It also houses the GPR and EMI electronics modules, the IR camera computer, the vehicle control processor, navigation system, telemetry modules, camera control modules and two lane marking systems.

The EMI detector is a Schiebel VAMIDS detector. DRDB specified its operating parameters and DRDB partially funded development of its precursor prototype. The sensor head array consists of 24 overlapping transmit/receive coaxial coil pairs mounted in a flexible plastic tray. The tray is dragged on the ground in front of the vehicle by a frame which projects from the front bumper of the vehicle. Maintaining sensor contact with the ground ensures that the target signal strength is maximized and the flexible tray allows the sensor array to ride over moderate sized obstacles. In-house designed algorithms process the signals from multiple coils to localize each detected object to within ± 10 cm across track and ± 5 cm along track.

Behind and above the EMI sensor array is a weather-proof container which houses an IR imager and a visible wavelength (VIS) imager. The IR imager is a modified Agema 1000 camera operating in a nominal 8-12 micron wave band and the VIS imager is a commercial CCD-based camera. The IR camera is equipped with a wide angle lens which has a field of view (FOV) of $62°$ across track by $41°$ along track. It is mounted at a height of roughly 2.5 m above ground and its optical axis is tilted approximately $40°$ below the horizontal. This yields a trapezoidal footprint on the ground whose near edge is close to the front edge of the EMI sensor array and is slightly greater than 3 m wide. The far edge is roughly 6 m in front of the near edge and is approximately 8 m wide. Images are

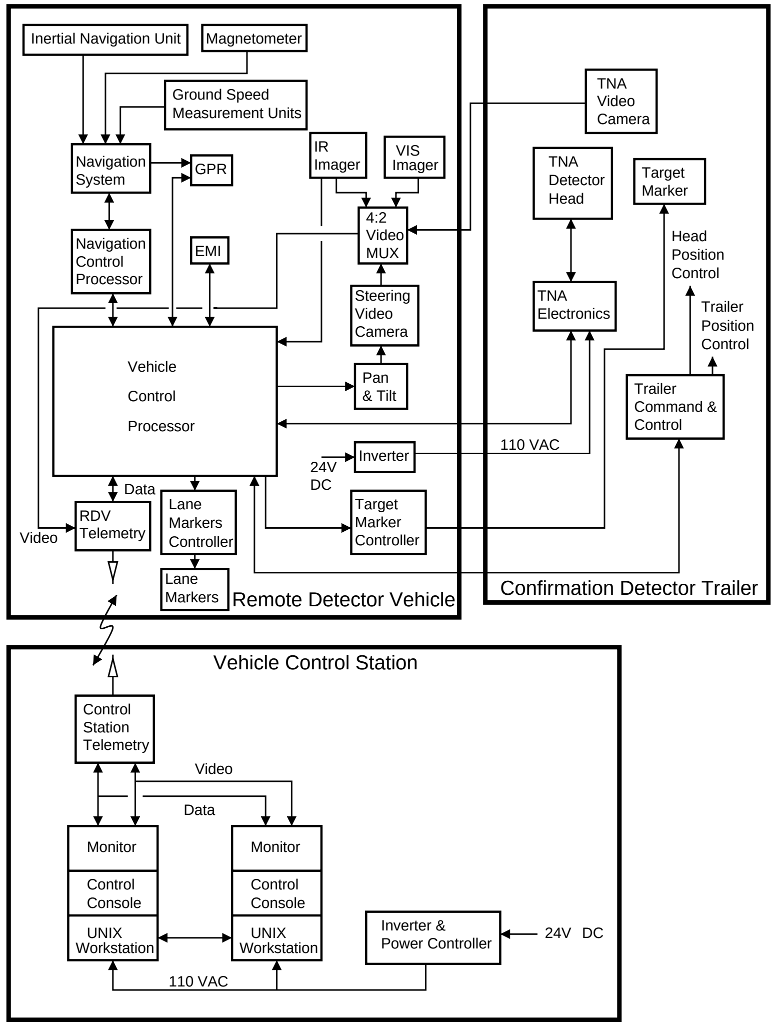

Figure 1. Functional block diagram of the major components of the ILDP system.

transmitted to the IR operator at 640 pixels across track by 400 pixels along track, which yields spatial resolution of less than a centimeter on the ground in the near half of the image. The FOV of the VIS camera has been set up to be roughly the same as that of the IR camera. Detection of mine-like targets is achieved by an operator monitoring the IR image at the vehicle control station. An automatic target detection (ATD) algorithm was deliberately not chosen because of the lack of robustness of available algorithms and a lack of time and funds to improve existing ones or to develop a new one. To improve probability of detection and reduce observer fatigue, dynamic range compression is done automatically prior to presenting the images to the operator. The imager has a large temperature range ($\sim \pm 200$C°) compared to the minimum sensitivity of the imager (< 0.1C°). This correponds to 4096 levels, whereas standard video can transmit a range of < 256 levels and the operator can discriminate < 64 levels. Dynamic range compression determines the portion of the range, which varies with time and conditions, that contains the scene temperature information within the region of interest and provides only that range for display to the operator. This is achieved using simple statistical-based compression methods. Contrast enhancement is currently done manually using the camera contrast controls, although automatic techniques are being considered for future implementations.

Behind the IR/VIS camera mount is the antenna of the GPR detector, manufactured by Elta Electronics. It is 3 m wide and is rigidly mounted 70 cm above the ground surface in a horizontal orientation in front of the vehicle.

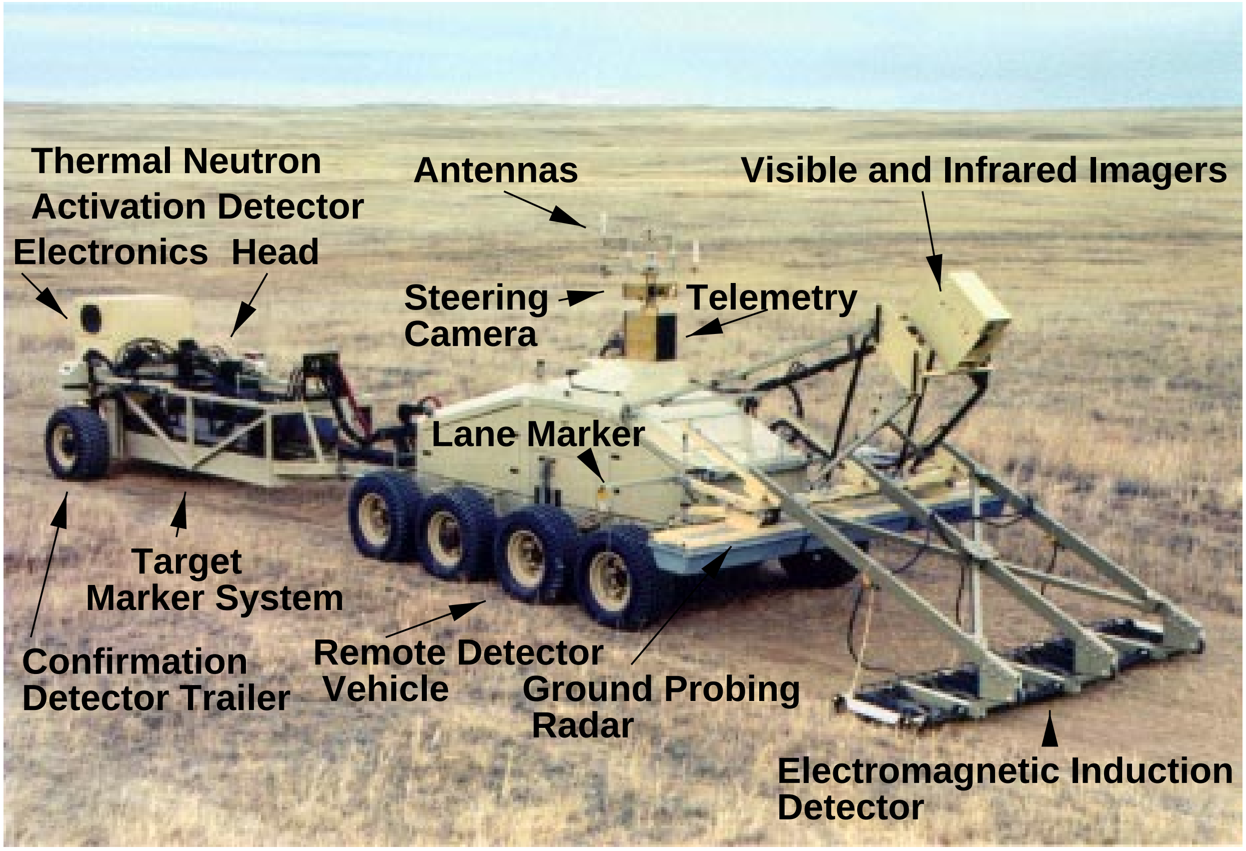

Figure 2. The ILDP system. The vehicle control station and its associated vehicle are not shown.

Its field of view is downward. The GPR provides as output a location of a target relative to its antenna position and a confidence level for the detection which is an indication of how "mine-like" the GPR perceives the target to be. The algorithms which produce these results are proprietary to Elta.

Command and control of the vehicle is handled remotely from a vehicle control station (VCS) mounted in a separate vehicle (not shown) a few hundred meters from the detector vehicle. High speed telemetry links transfer command and control information to a vehicle control processor on the RDV. The vehicle control processor communicates with the detectors, navigation system, and other subsystems on the vehicle. Command and control information includes data necessary for driving the vehicle, video signals from the steering cameras or camera mounted on the TNA detector, command and control information for all detectors, detector data, IR and VIS camera video and C-code Global Positioning System (GPS) data (for logging detections). The data fusion, the EMI and TNA detection algorithms and the vehicle command and control are processed by the VCS.

Lane markers mounted on the sides of the vehicle provide a continuous indication of where the vehicle has been. Both the lane marking and the target marking systems use a novel apparatus that mixes a dry powder with a quantity of dyed water immediately prior to marking. The resultant mark consists of viscous, fluorescent jelly-like material that is highly visible at up to 20 m day or night, but is biodegradable and disappears after a day or two. The ability to disappear is a requirement to allow follow up clearance within a few hours, but still maintain the ability to operate the detector system on the same road a day or two later without previous marks causing confusion. Lane and target markings can be made different colours. Since the ratio of powder to water is very small, very little powder is needed for a mission and sufficient water and powder for an 8 hour shift is easily stored on the vehicle.

Mounted on a trailer behind the detector vehicle is a TNA detector which is used to confirm the presence of a mine at a suspect location. The detector consists of two parts, a head which sits on a moveable carriage and signal processing electronics with a controlling computer. A target marker nozzle is mounted coaxially with the center of the TNA detector head. The TNA electronics is situated in a climate controlled enclosure at the back of the trailer. The TNA head is first coarsely positioned by moving the entire trailer along-track with respect to the RDV by means of a hydraulic piston which couples the two platforms. Next the trailer wheels rotate 90° and the trailer moves in an arc to the approximate desired location. Finally, the TNA head is finely positioned by moving the carriage in which it sits along-track to the desired position. The head is then lowered almost to the ground surface and a reading is taken for a fixed time period. The entire process of moving the TNA head and taking a reading may be done either manually or automatically under computer control. The confirmation detection is unique to the ILDP system and is essential to reduce the combined false alarm rate of the scanning sensors to an acceptably low rate for the system.

2.3. Description of Operation

Two operators control the system from the vehicle control station. Each sits in front of one of two almost identical workstations. Although the displays and controls on one workstation appear different than those of the other, they can be swapped by software commands to enable each to behave like the other. One console displays the image from the steering camera, various vehicle controls and indicators. It also has a scrolling data fusion screen. The operator at this console drives the vehicle using a joy stick, track ball and several switches. He/she also monitors the data fusion screen in case operator intervention is required.

The second console displays IR imagery and its operator is responsible for detecting mine-like anomalies in the moving imagery. When a target of interest is discovered, the operator moves a cursor on the target using a track ball and presses a button to designate the target. Although the target may be viewed and designated by the operator as it passes through the entire image, designation is typically done in the bottom third of the image where error in detection location is minimized due to the reduced error in using the flat earth approximation and improved spatial resolution. Also detectability is enhanced due to the increased perspective angle and larger target size. Additional buttons may be pressed to coarsely classify the target as strong/weak or big/small. The operator can also switch from the IR to the VIS imager if a particular target appears confusing. This is normally only done after halting the vehicle. Once a target is indicated, it is automatically tracked by a custom designed target tracking algorithm until it gets near the bottom edge of the IR FOV, where the frame buffer coordinates are read. These coordinates are transformed to the vehicle reference frame and then, using navigational information, to a local world-based coordinate system called DRABS (Dead Reckoning Absolute). The confidence of the target (strong/weak or big/small), together with its DRABS coordinates, are passed to the spatial correspondence software in the vehicle control station. At the same time, the spatially registered detection is displayed on the data fusion screen as an icon (Fig. 3) which scrolls from top to bottom as the vehicle moves.

As the vehicle advances, the area which was scanned by the IR imager passes into the field of view of the EMI detector. The time sequence of responses from each sensor head pair together with the vehicle motion produce an intensity map with pixel size a function of coil width and vehicle speed. The coordinates of each pixel are transformed to DRABS coodinates and the coordinates and associated intensity are passed to the VCS. There an automatic target detection algorithm examines the intensity map and declares detections, some of which correspond to targets indicated by the IR imager and some of which do not. The DRABS coordinates of each EMI detection and its associated confidence (a weight or "strength"), are passed to the spatial correspondence software. The spatially registered detections are also displayed on the data fusion screen as icons whose shape and colour differ from those of the IR imager. The raw EMI signal is also displayed to the operator as a background to the data fusion display.

The area which was scanned by the EMI detector next passes into the field of view of the GPR detector. The detailed functioning of the GPR is proprietary, but it is speculated that it produces an internal return signal intensity map, somewhat analgous to the EMI detector. This map is then analysed to extract features, based on local shape and intensity information, which are then used to classify the returns as being a mine with a specified degree of confidence. The GPR detection algorithm, internal to its own processor, uses navigational information provided to it by the ground speed measurement units. For each detection, the GPR outputs a set of coordinates in the detector coordinate system and a confidence level. The coordinates are converted to DRABS coordinates and are passed, with the confidence level, to the spatial correspondence software and the data fusion screen for display as an icon.

Each detector has a lag time between its instantaneous FOV passing over a target and the target being registered as a detection. Once this has compensated for, the spatial correspondence software can determine if different detections correspond to the same target. The results of the spatial correspondence software are passed to the scanning sensor fusion software which decides if the resulting target is sufficiently significant to be considered a suspect. If a suspect target is identified, an additional icon is placed at that location on the data fusion screen. The driver operator can override the system at any time by deselecting a suspect target. However, if this has not been done prior to the suspect reaching the confirmation detector workspace boundary, the vehicle is automatically brought to a controlled stop.

Once a suspect target enters the confirmation detector workspace and the vehicle has halted, the TNA head is automatically positioned over the target location, which has been estimated by the data fusion/spatial correspondence algorithm. If desired, the operator can view the positioning of the head using a TNA camera which inspects the TNA head field of view. If surface clues give additional information to indicate that the estimated position of the suspect target is incorrect, the TNA positioning can be completed manually. Once the head is positioned, lowered, a measurement taken and the head raised again, the final data fusion algorithm indicates to the operator its confidence in whether a mine is present. The operator then decides whether to trigger the target marker, in which case a quantity of marking material is ejected underneath the center of the TNA head. Detections can be logged for reporting purposes using a C-code global positioning system (GPS).

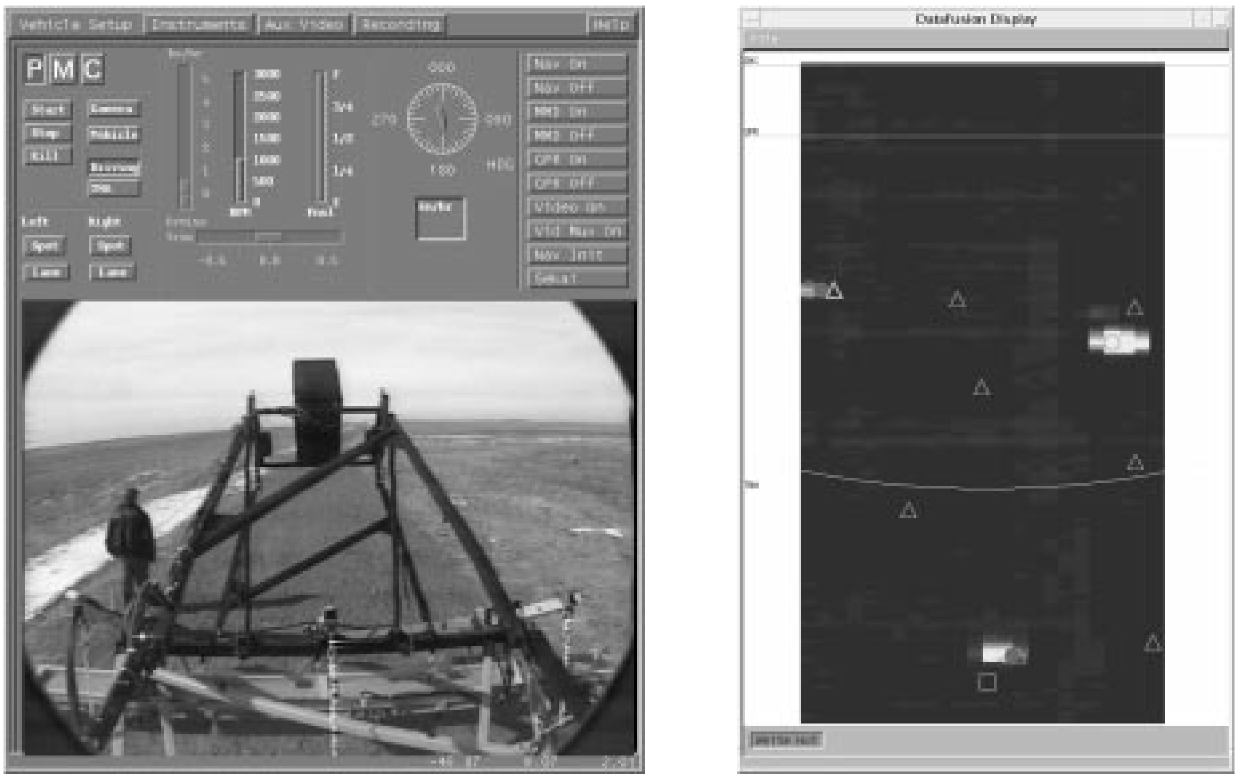

Figure 3. The driver's console display. The left hand portion of the screen contains system controls and indicators at the top and the view from the steering camera below. The data fusion screen is at the right. A gray scale map from the EMI detector moves from top to bottom as the vehicle moves. Hollow icons represent detections by the individual detectors (IR - squares, EMI - circles, GPR - triangles). Icons with thicker lines are associated with higher confidence detections for each detector. The solid icons are detections by the scanning sensor fusion algorithm. The horizontal lines correspond to the rear edge of the FOV of the EMI and GPR detectors. The curved line is the forward edge of the TNA workspace. The IR operator's console display is not shown.

The method of determining whether detections from different detectors correspond to the same target and hence can be combined and the method of combining them is critical to the success of the system. It is discussed in detail in the next section.

3. SPATIAL CORRESPONDENCE AND DATA FUSION ALGORITHM

3.1. Overview

Each of the scanning sensors (IR, EMI, GPR) provides information concerning the presence (or absence) of physical properties which accompany the presence of landmines. For example, IR provides a measure of thermal anomalies, EMI reports anomalies in electrical conductivity, and GPR detects anomalies in dielectric and other electromagnetic properties. As the system operates, each scanning sensor observes a different area of the path according to that sensor's position on the RDV and its own field of view. A detection from any sensor consists of a position estimate and a measure of the confidence that the respective property, detectable by that sensor, is present in a local patch of ground about the reported position. For the IR and GPR, the positional information (which may include a time lag) is with respect to the sensor reference frame. Due to the image-based nature of the EMI detection algorithm, its positional information is in the local world-based coordinate system (DRABS).

Temporal and spatial delays are unavoidable in the detection positional information. For example, transmission delays from the RDV to the control station are temporal, and where possible are managed by time stamps on the RDV. Target detection algorithms for the EMI and GPR result in time delays proportional to vehicle speed since minimum distances must be travelled before sufficient information is available to declare a detection.

Before detection-level data fusion can be applied, all detections from the scanning system must be registered spatially in the same frame of reference. This requirement demands that highly accurate geometrical parameters and estimates of vehicle motion be available. Geometrical information is computed through geometrical calibration while estimates of vehicle motion are provided by the navigation system.

Once all detections are spatially registered in a common frame of reference, spatial correspondence algorithms can be applied. Spatial correspondences are decisions derived by the system which relate individual detection events according to whether or not these events originated within the same local patch of ground, and hence could have originated from a single landmine. If the overall confidence level for detections determined to be in correspondence is significant, a position for placement of the TNA confirmatory sensor is computed, and the system gradually decelerates to a stop with the suspect patch of ground within the TNA work space.

Information provided by the TNA following deployment over a local ground patch is used in combination with scanning sensor information from a corresponding patch in a final fusion stage. This process results in an overall confidence level concerning the presence (or absence) of a landmine at that patch. A significant confidence level at this stage results in the system recommending that the target marker be triggered, although the operator makes the final decision to trigger the marker. All relevant information concerning that detection (or non-detection) is then recorded.

The overall data fusion process, then, consists of six main components; geometric calibration, the navigation processing, spatial registration, spatial correspondence, scanning sensor fusion and confirmation fusion. Geometric calibration is performed off-line prior to a mine detection mission. Navigation processing is handled by a separate navigation computer on board the RDV. Spatial registration, spatial correspondence, scanning sensor fusion and confirmation fusion computations are all performed on the vehicle control station. Each of these components is described further in the following sections.

3.2. Geometric calibration

Geometric calibration refers to the process used to estimate the parameters required for reference frame transformation and optical and auxiliary sensor calibration. Each primary detection sensor (IR, EMI, GPR, TNA) has its own frame of reference. So too does the vehicle, the navigation system and its components, and all auxiliary encoders and sensors which measure relative positions or angles of system components. Geometrical calibration is performed through physical measurement and data collection in a controlled environment, followed by post processing. The overall process results in numerical parameters for translations and rotations relating the various reference frames to one another. This information is essential in order to transform detection positional information, originally reported relative to a sensor reference frame, to the vehicle reference frame. Optical calibration of the IR subsystem is also performed so that operator designations within the displayed imagery can be transformed to positional vectors relative to the vehicle reference frame. Proper calibration of auxiliary sensors which are used to determine the TNA detector position relative to the vehicle is also essential.

3.3. Navigation processing

The navigation problem is one of state estimation which filters and transforms raw navigation sensor information to derive robust and highly accurate estimates of the motion state of the system. The vehicle motion state consists of translational velocity, translational acceleration, attitude (roll, pitch, heading), and angular velocity. Navigation sensors include encoder wheels for ground speed measurement, a three-axis accelerometer unit, a three-axis rate gyroscope unit, and a magnetometer. Differential GPS (DGPS) measurements of vehicle position can provide sufficient accuracy for data fusion, but its use is limited if the system is intended for other than R&D purposes. Reliability problems in some geographical locations, including heavily treed or urban areas, and the logistics considerations of setting up a base station make it awkward and sometimes impractical for a widely used operational system. Thus, DGPS is not currently being used in the estimation of the vehicle motion state, although C-code GPS measurements can be made available for reporting purposes. Measurements by the navigation sensors are provided as inputs to the navigation processor which derives the motion state through the use of Kalman filtering. The navigation processor resides on the RDV, and the derived motion state is transmitted to the control vehicle for use within the data fusion algorithms. A short navigation history is maintained within the VCS software so that spatial and temporal delays from the sensors can be corrected before transforming the detections to a common reference frame.

3.4. Spatial registration

Spatial registration is the process of transforming detection information from any of the scanning sensors to a common frame of reference. The accuracy of this process is highly dependent on the accuracy of geometrical calibration and to a lesser degree on the navigation filters. For example, an IR detection is initiated by the operator designating a position within the displayed image on the vehicle control station. This position is automatically tracked in image coordinate space to the trailing edge of the image and then the position (frame buffer row and column numbers) is transformed using optical parameters of the geometrical calibration process to derive a direction vector from the IR reference frame to the suspected point on the ground in the IR field of view. A local flat-earth approximation is used to compute the position of the detection relative to the IR reference frame. This position is transformed to the vehicle reference frame using geometrical calibration parameters, and is further transformed to the common DRABS reference frame using navigation and temporal lag information. Similar processes are used for EMI and GPR detections, but with additional corrections for spatial lag in the detection information. Once detections are spatially registered in the DRABS reference frame, the detection information can be displayed to the operator(s) in a spatially correct manner. Spatial registration within a common reference frame is essential before spatial relationships between detections can be inferred.

3.5. Spatial correspondence

The output of the spatial registration is a set of detection events, each with a location estimate in a common reference frame and an uncertainty in the location estimate. The spatial correspondence algorithm is then applied to partition the set of detection events into classes, with each class representing those detections which could have resulted from the same local patch of ground or a single target. The correspondence decision for any two detections takes into account the relative location of the detections and the corresponding localization uncertainties. The spatial extent allowed for correspondence is dictated largely by the spatial extent of the TNA field of view.

3.6. Scanning sensor fusion

The information contained in detections within a correspondence class is used to determine an overall position and confidence level for the suspect patch of ground. The overall confidence level is derived through a weighted voting strategy in which the weights are computed based on environmental and operational parameters. For example, poor IR conditions or dense metal clutter in the operational area can be used to adjust the corresponding weights assigned to IR and EMI detections. Estimation of weights for a detector, given a particular response, requires *a priori* knowledge of signal properties for a set of expected targets with respect to those of typical *in situ* background signals, for given environmental conditions and geographical location. This can be determined before commencing the actual mine detection operation by several methods. Background signals can be monitored for each sensor and set of ambient conditions appropriate to the sensor, such as soil type for an EMI, GPR or IR sensor, or diurnal air temperature profiles for an IR sensor. Measurements of unarmed mines or surrogates, deliberately buried in a test section of road, can be made prior to starting the actual mine detection operation. Time varying parameters which are well correlated with detector performance, but are not simply related to environmental parameters, can also be monitored. Examples include measurement of soil temperature gradient profiles, either locally or at a remote site, for IR detection[3] and measurement of soil moisture content for all scanning sensors and the TNA. If the overall confidence level for a correspondence class is significant, a position for placement of the TNA confirmatory sensor is computed and the system is brought to a halt with the suspect location within the TNA workspace.

At present, a simplified version of the scanning sensor fusion is employed, which has been labeled "2+ fusion". A detection is declared if any three sensors declare a detection with a low confidence level or any two declare a detection with a medium confidence level or any one detector declares a detection with a high confidence level. The confidence levels are determined in advance based on precalibration in the local environment. Locational errors are assumed to be the same for all sensors and are again determined in advance based on calibration results. The simplified version has been used for preliminary data analyses and as a benchmark during the development of the full scanning sensor fusion algorithms.

3.7. Confirmation fusion

Measurements from the TNA system generate a confidence level that the local patch of ground under observation contains a sufficient amount of nitrogen to indicate the presence of a landmine. This confidence level is combined with the scanning system confidence level for a corresponding patch (recognizing that the patches will not generally coincide due to navigational errors) to generate the system confidence that this patch contains a landmine. If this system confidence level is significant, a detection is recommended. The operator makes a decision to trigger the marking system based on the recommendation and additional strategic or tactical knowledge that might be available. If a detection is declared, all data relevant to this detection event is recorded by the VCS.

4. SYSTEM TRIALS

Testing and evaluation of individual sensors has been ongoing since Winter 1994 and development and testing of data fusion methodology has been underway since 1996 using a non-teleoperated surrogate vehicle instrumented similar to the ILDP system. Results of individual sensor experiments have been reported previously[2,4] and are also reported in these proceedings.[5]

The first full system trial was conducted between November 12 and 28 1997. The aim of the trial was to provide a fairly realistic but tough detection scenario, which approximated operational conditions. Mines were buried at the Defence Research Establishment Suffield in Dugway Trail, a well compacted dirt road approximately 5 km long. The trial comprised 11 active detection days with an expenditure of roughly 4 person-months. During 32 hours of actual detection time, 78.5 km of road were covered, comprising an area of 235,000 m^2, which yielded an average speed of 2.4 km/h. In total 759 mines were traversed, of which 249 (32.8%) were metallic and 510 (67.2%) were low metal content. Mine types, which were all AT, consisted of low metal Type 1 (26%), low metal Type 2 (8%), low metal Type 3 (7%), low metal Type 4 (26%), metal Type 1 (6%), metal Type 2 (23%) and metal Type 3 (4%). Mines were defuzed but an amount of metal equivalent to that in the fuze was placed in the fuze wells of the low metal content mines. Mines were buried using tactical methods between 3.8 cm and 17.8 cm depth (top of mine to ground surface), with an average depth of 10.2 cm. Ground truthing of mine positions was conducted by surveying at burial time using a 2 cm accuracy real-time DGPS. The trials were "blind", that is, except for a calibration section at one end of the road, the operators were not told the locations, numbers or types of the mines and whether or not potential false targets, such as refilled holes, would be present.

Night and day operations were conducted and both temperate and cold weather conditions were encountered. The system functioned well under most of the conditions, although flat diurnal temperature profiles, fog and ground frost led to suboptimal IR performance for many of the trials. A few experienced operators were used together with a large number of neophyte operators. Although data fusion software was available, real-time data fusion was not implemented on the VCS at the time of the trial. Real-time electronic or physical marking were also not done. Instead, detection, vehicle navigation, and DGPS data were simultaneously acquired and processed off-line. The DGPS data were required for ground truth comparisons, but were not used in the vehicle navigation processing. The simplified data fusion method, described at the end of the "Scanning sensor fusion" subsection, was employed for the analyses that follow. Finally, the TNA confirmatory detector was not employed for these trials since the automatic trailer control was not yet implemented.

A number of variables affect the measured system performance. Some can be controlled for a given experiment, such as target type and depth, detection algorithms, sensor fusion methodology, and sensor performance reproducability. Less controllable are equipment condition, particularly in the "infant mortality" stage of a new system and operator experience and aptitude, which can dramatically affect results. Scoring criteria can be fixed for a set of experiments, but a detailed knowledge of the criteria is critical if a meaningful comparison of two sets of experiments scored according to different criteria is to be made. Mines that are outside the detection width of the vehicle can be rejected from analysis. However, it is difficult to ensure that a detector passes over a target the same way each time and this variable will affect the measured performance of that detector. Finally, one has little control over environmental conditions which can seriously effect the performance of various sensors. The result of this variability is that there can be no single number that adequately describes a system's performance.

In order to score results, it is necessary to define when a detector alarm is considered to be a detection of a true target. The conventional manner is to define a distance or "halo" around the target center or edge (the center was used for these trials) and any alarm within that distance is considered a detection. The selection of this distance is difficult to determine. Ideally, the halo should be a different size for each detector, based on its location estimate uncertainty for each target. The latter may vary with target type or location or with environmental conditions. Because the full data fusion module was not complete and in particular accurate target location estimate uncertainties were not available at the time of the trials, a simpler approach was used. Estimated probability of detection (P_d), corrected for random false positives within halo, was plotted versus halo radius for a subset of the data. As expected, P_d increases initially with halo radius as more true detections are included and then starts to decrease as no more true detections are included and the false positive correction dominates. The position of the peak of the P_d curve, which occurred in the vicinity of 60 cm for all three detectors, can be considered the optimum halo radius. This is also consistent with estimates of the positional accuracy and location uncertainties which were between 40 and 60 cm at the time of the trial. Thus, 60 cm was chosen as the halo radius for the remainder of the trials.

The estimated probability of detection and estimated false alarm rate (FAR) for all runs of the November 1997 trials are shown in chronological order from left to right in Fig. 4. A wide variation in values is seen, due to the variables discussed above. An additional source of variation is that for some runs, very few mines were passed over. In such cases, the small sample size will yield a large uncertainty in the estimate of the probability of detection. Note the gradual increasing trend in P_d from the first run up to roughly the middle. This is because at the beginning of the trial most of the operators were inexperienced, as was the support crew. As time progressed, fewer inexperienced operators were used, the support crew gained experience and became familiar with the system and a number of bugs were removed from the system.

Figs. 5 and 6 show that for metal AT mines there is little variation of P_d with mine depth, over the range involved, for all individual sensors and for the data fusion output. The same is true for low metal content AT mines for the IR imager. However, there is a decrease of P_d with depth for low metal content AT mines when the EMI and GPR detectors are considered. This is because the metal mines provide a strong signal sufficiently above background for the EMI and GPR detectors at all depths. The same is true for the IR imager signal from AT mines of either type.

If we assume that in the last half of the set of runs, crew inexperience no longer had a negative effect on results, an average of these runs will be a measure of the system performance as effected by all uncontrollable variables. The mean and standard deviation of the estimated probabilities of detection and false alarm rates, averaged over runs 31 to 75, are shown in Table 1.

As previously mentioned, the above results do not include confirmation using the TNA detector. Based on independent experiments,[2,4] its use is expected to yield a slightly lower ($\sim 3-5\%$) overall P_d but reduce the FAR by more than a factor of ten.

5. CONCLUSIONS

The ILDP system is a teleoperated vehicle-mounted multisensor system for the detection of landmines on roads and tracks in rear area combat situations and Operations Other Than War. The system is easy to use. It has a number of features that are novel for such systems. It is a fully integrated and operational system, rather than just an R&D platform. It employs an integral real-time data fusion system to reduce false alarm rates and it employs a confirmatory detector, TNA, to further reduce the FAR of the initial scanning detectors. In part because of its easy to use displays, moderately skilled operators can be trained to use the system effectively in about one week.

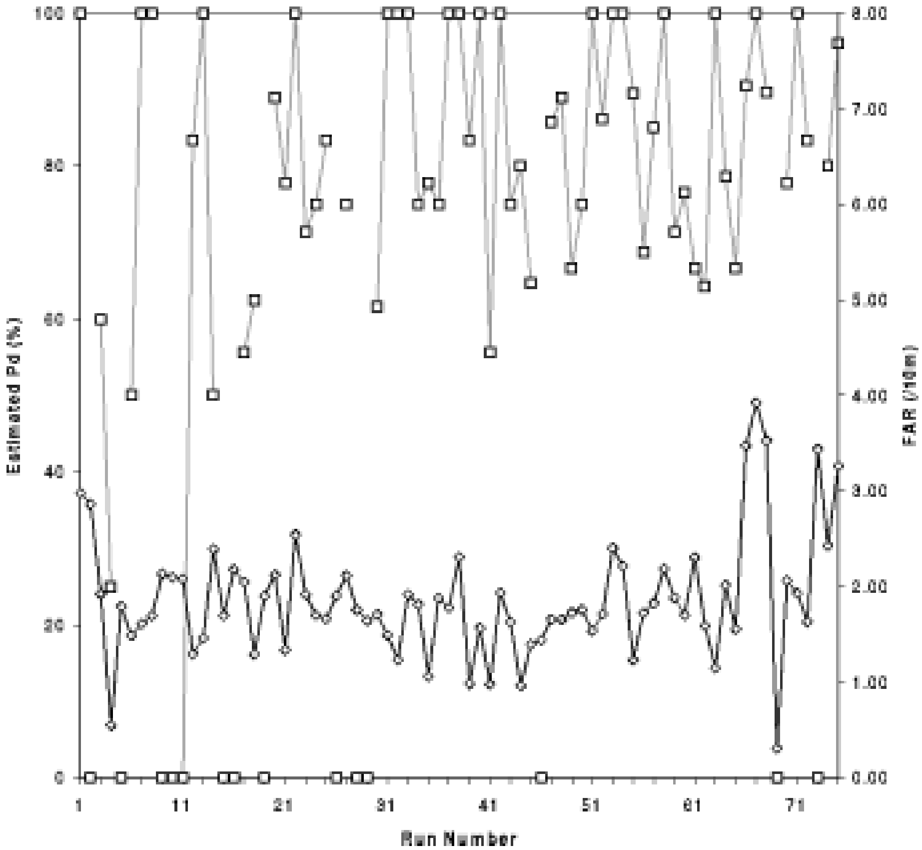

Figure 4. Estimated probability of detection (squares) and false alarm rate (diamonds) for all runs of the November 1997 trials. The "2+ fusion" method described above has been used. Run numbers are in chronological order. Halo radius is 60 cm. With the exception of run 11, all P_d values of 0 denote that no mines were passed over in the run.

Target Class	P_d (%)		$FA/10$m	
	Mean	Standard Deviation	Mean	Standard Deviation
All mines	78.31	14.77	2.01	0.73
AT Mines	85.08	13.25	2.01	0.73
Metal AT mines	100.00	0.00	2.22	0.73
Low metal AT mines	77.84	20.72	2.01	0.73
AP mines	37.30	32.22	2.72	0.73

Table 1. Summary of Results of Second Half of November 1997 Mine Detection Trials

Future effort will be aimed at using the lessons learned from the trials to improve the general reliability of the system. The P_d and FAR can be improved by reducing navigational errors due to the navigation sensors and processing algorithms. Completing and implementing the full data fusion algorithm should also improve P_d and FAR. Improvements to existing detectors and replacement by better detectors of the same technology or different technologies will also be investigated. One important result of the trials is the realization that proper operator selection and training is critical and ways to improve both will be investigated.

The ILDP system is scheduled for testing at the US Governments's GSTAMIDS Advanced Technology Demonstrator trials in Summer 1998. A follow-on project has been initiated to build four to six production units for the Canadian Forces by the year 2000 and the contract is expected to be awarded shortly.

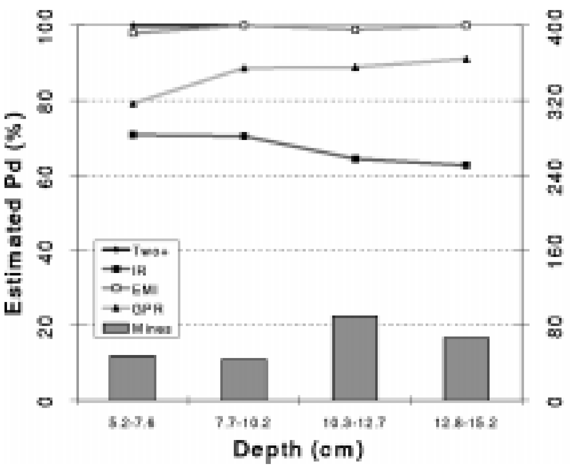

Figure 5. Estimated probability of detection versus mine depth for metal mines. Halo radius is 60 cm. "2+" refers to the data fusion method described above. Number of mines used for each estimate is shown on the bar graph.

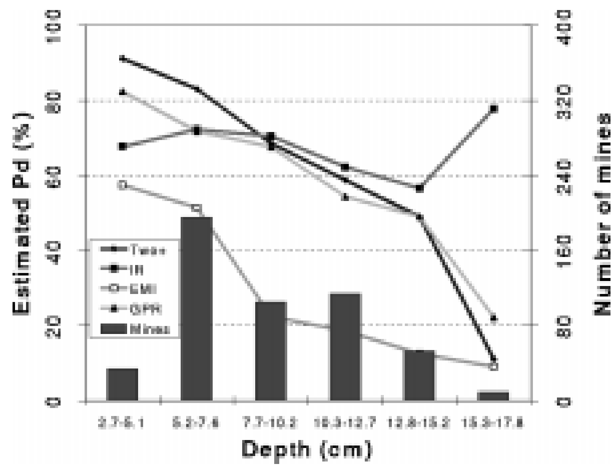

Figure 6. Estimated probability of detection versus mine depth for low metal content mines. Halo radius is 60 cm. "2+" refers to the data fusion method described above. Number of mines used for each estimate is shown on the bar graph.

6. ACKNOWLEDGMENTS

We wish to acknowledge the many DRDB team members who contributed substantially to the design, construction and testing of the ILDP system, including; Doug Benson (GPR and Telemetry), Chris Brosinsky (Vehicle Mechanics), Stan Burton (Vehicle Control), Major Al Carruthers (Project Manager), Tom Cousins (TNA), Ron Eirich (Vehicle Control), Rob Ellingson (Marking System and Vehicle Mechanics), Capt. Mark Hache (Trials Coordination), Doug Hanna (Vehicle), Bob Herring (Navigation), Hank Hudema (Telemetry and Trials), Master Gunner (Ret.) Bill MacLean (Trials), Terry Meidinger (Systems Integration), Steve Penzes (Vehicle), Jean-Robert Simard (IR), Wayne Sirovyak (IR), Jim Sodero (Vehicle Control), Jack Toews (EMI), Benny Wong (GPR). We also wish to acknowledge the efforts and cooperation of the many contractors who worked closely with the team throughout the project.

REFERENCES

1. J.E.McFee, Y.Das, A.Carruthers, S.Murray, P.Gallagher, and G.Briosi, "CRAD countermine R&D study - Final Report (U)," Report SSP 174, Defence Research Establishment Suffield, April 1994.
2. J.McFee and A.Carruthers, "A multisensor mine detector for peacekeeping - Improved Landmine Detector Concept (ILDC)," in *Detection and Remediation Technologies for Mines and Mine-like Targets*, A.C.Dubey, R.L.Barnard, C.J.Lowe, and J.E.McFee, eds., *Proc. SPIE* **2765**, pp. 233–248, 1996.
3. K.Russell, J.McFee, and W.Sirovyak, "Remote performance prediction for infrared imaging of buried mines," in *Detection and Remediation Technologies for Mines and Mine-like Targets II*, A.C.Dubey and R.L.Barnard, eds., *Proc. SPIE* **3079**, pp. 762–769, 1997.
4. T.Cousins, T.A.Jones, J.R.Brisson, J.E.McFee, T.J.Jamieson, E.J.Waller, F.J.LeMay, H.Ing, C.E.Clifford, and B.Selkirk, "The development of a thermal neutron activation (TNA) system as a confirmatory nonmetallic land mine detector," *Journal of Radioanalytical and Nuclear Chemistry* , January 1998.
5. J.McFee, T.Cousins, T.Jones, J.R.Brisson, T.Jamieson, E.Waller, F.LeMay, H.Ing, E.Clifford, and B.Selkirk, "A thermal neutron activation system for confirmatory non-metallic land mine detection," in *Detection and Remediation Technologies for Mines and Mine-like Targets III*, A.C.Dubey, J.F.Harvey, and J.T.Broach, eds., *Proc. SPIE* **3392**, 1998.