Casi Real-Time Surface-Laid Mine Detection System

Tyler Ivanco, Steve Achal, John McFee*, Cliff Anger

Itres Research Limited
Calgary AB
(403) 250-9944 (Vox)
(403) 250-9916 (Fax)
[tivanco|sachal|canger]@itres.com

*Defence Research Establishment Suffield (DRES)
Box 4000
Medicine Hat, AB CANADA T1A 8K6
(403)544-4739 (Vox)
(403) 544-4704 (Fax)
John.McFee@dres.dnd.ca

*A ground vehicle-based, real-time, surface mine detection system, utilizing a Compact Airborne Spectrographic Image (*casi*), efficient mine detection algorithms, and real-time processing systems, was designed and tested. The combined real-time system was capable of "learning" the in-situ spectra of various mines, thus providing a spectral library for the detection algorithms. The real-time processing of the casi data involved three steps. The first step was the radiometric correction of the raw data. The second step involved the application of the mine detection algorithms to the corrected data, referencing the spectral library. In the final step, the results of the real-time processes were stored and displayed, usually within a few frame times of the data acquisition. To the authors knowledge, this system represents the first hyperspectral imager to detect mines in real-time.*

This paper describes the generation of the in-situ mine spectral library, the collection of the scene data, the real-time processing of the scene data and the subsequent display and recording of the detection data. The limitations and expansion capabilities of the real-time system are discussed as well as various techniques that were implemented to achieve the goals. Planned future improvements that have been identified to create a more robust and higher performance, yet simpler processing system, are also discussed.

1. Introduction

Itres and Defence Research Establishment Suffield (DRES) have been involved in a long-term cooperative research study on the detection of surface laid mines of a variety of types. Past research has focused upon analysis of the spectral reflectances from various mine types under various environmental conditions and the requisite algorithms necessary to detect these targets. The practical issues of vehicle platforms (e.g. aircraft, helicopter and ground vehicles), instrument throughput (directly impacting spatial resolution) and environmental variances were also studied.

Data from the *casi*, like that from other imagers which are intended for remote sensing, is stored during a mission and analyzed off-line after the mission. While this is reasonable for remote sensing applications, including humanitarian demining, it is not acceptable for the military in a conflict situation, who must be able to identify a minefield within a few minutes of flying over it. Thus, recent R&D has been aimed at demonstrating the feasibility of detecting mines in real-time. The goal is to develop a system that can detect mines from an airborne platform at typical speeds. To reduce risk, time and funds, however, it was decided to initially aim for a speed improvement that would allow real-time detection from a slow, land-based vehicle. This would require a minimum of modifications to the *casi* itself, but would lead to a choice of some design elements that would not be extendable to further increase speed later on. Such an instrument could still be very useful as an additional or replacement sensor on a teleoperated vehicle-mounted multi-sensor mine detector system such as ILDP (McFee *et al*, 1998) or on the protection vehicle that precedes it.

The result of this recent research is a real-time processing implementation that utilizes the *casi-2* instrument and the algorithms previously developed to detect the targets. The processing platform involves a system that generates the results of the data acquisition and target analysis to an operator by displaying probability information alongside the base imagery. The entire process completes within a few frame times of acquisition (a frame time is approximately 15-35 ms.).

The ability to acquire, correct and analyze *casi* images in real-time will, as well, provide significant benefits for its traditional remote sensing users. The project has led to a commercial prototype for use in Itres airborne data acquisition systems. The advantage of such a real-time processing system to acquire, radiometrically correct, and "fuse" data from a variety of systems leads to higher quality control of all data streams and ultimately lower operating costs. This capability permits the operator to note and possibly resolve problems in the data stream prior to leaving the area of interest saving time and reducing mission costs necessitated by re-flights or data processing failures.

Detection of the targets requires two modes of operation from the system. In the first mode, the system learns ("sniffs") the target spectra storing the profile for later use. During the identification stage, the system searches out targets ("fetches") by acquiring spectral data for each pixel then applying the comparative algorithms to the candidate pixel, using the stored reference spectra. The comparison may be applied a number of times for each of the possible sets stored during the learn mode. The results are then compiled for display along with the input imagery and if necessary, any ancillary information (such as location) that may be collected.

2. General Description of the Solution

The *casi* collects frames of spectral data as defined by the instrument setup. This data must be processed to correct for various instrument characteristics in order to generate radiance values. These values are subsequently compared against reference spectra to determine whether a candidate pixel conforms to known target spectra. This process needs to operate in near real-time (i.e. less than 0.5 second response) and minimize false positives. It should be self correcting (i.e. change reference spectra to match environmental conditions) and automated, requiring little if any intervention during acquisition.

Conventional *casi* processing takes the form of:

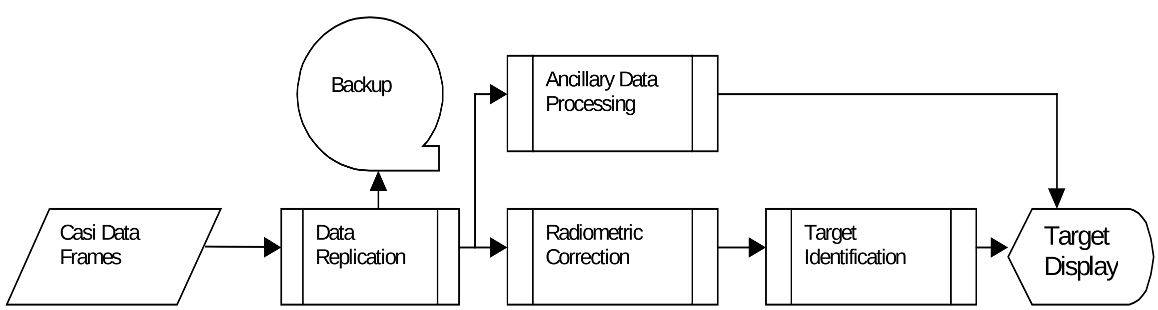

- acquire image data onto mass media (e.g. tape or disk)
- copy the image data to a backup media
- process the image data through a radiometric correction module creating a data base which is compatible with the PCI remote sensing and imaging software[1]
- process ancillary data (such as from a Global Positioning System (GPS) and inertial measurement unit (IMU)) for geo-location functions
- apply analysis algorithms to either learn ("sniff") or identify ("fetch") using a target library augmenting the PCI database image information with results of the analysis
- display the results and apply threshold limiting steps to minimize false positives and adjust sensitivity minimizing the number of missed targets.

The previously developed algorithms operated within the EASI/PACE PCI processing environment without any regard to real-time issues. The nature of this process permitted an analyst to tailor the algorithm's parameters, maximize the sensitivity and minimize false alarms interactively. It was necessary to reduce this data processing complexity so that the entire process could be automated.

The large volume of data generated by this and possibly future instruments requires a scalable system. In this context a scalable system is one where additional processors would have a dramatic improvement (ideally linearly proportional to the number of processors) on the productivity of the system. Scalability permits the development of not only higher throughput systems, but also systems with more complex processing requirements.

The current algorithms depend only on the spectral characteristics of the pixel; this achieves a frame independent processing stream. The action of pipelining the frame imagery permits a form of coarse parallelism. The backup process can be performed as data is passed onto the radiometric process. The radiometric and target identification processes can be applied independently to each frame; that is, the results of processing a frame will not affect the results of processing other frames. In this type of situation, it is easy to replicate the radiometric correction and target identification processes to a number of frames with the results passed back to a collation process that reorders the frame results for later display.

A PCI bus system, with embedded microprocessors based upon the Intel technology, was chosen. This technology is powerful, compact, inexpensive and readily available. A variety of software, from operating systems to development tools, is available with strong upgrade potential.

Each real time system was deployed within a rack mount system with redundant power supplies and four industrial processor systems all running autonomously. Additional units could be replicated providing a simple mechanism to scale up the total processing capability.

The identical nature of each processing unit permitted a great deal of flexibility. The processing actions could be assigned to any processor or any number of processors by changing the software installations.

3. Target Detection System Requirements

Extracting information from the acquired spectral data requires a significant amount of processing. The processing system must be capable of handling software that controls the basic imager which detects spectral differences, identifies those differences and displays the results to an operator in real-time.

---

[1] PCI Geomatics, www.pcigeomatics.com

The casi is a flexible visible and near infrared instrument with a spectral range of approximately 400nm to 950 nm and a spectral resolution of approximate 2 nm. The number of bands, location and bandwidth are user-definable. The instrument collects up to 512 spatial pixels at the user-defined bands at frame rates exceeding 100Hz. These frame rates are dependant upon the number of bands collected and the general configuration of the system. Overall, the system is currently capable of generating spectral information approaching one Megapixel per second. Future enhancements anticipate two to three fold improvement in the data generation rate. New instruments in development have a potential ten-fold increase in data acquisition speed.

The *casi* instrument transfers data from the sensor array to an on-board SCSI controller and display in real-time with minimal data manipulation. However, prior to storage, this data is collated with ancillary information. These storage units are either Mammoth 8900 Exabyte tape systems or an array of removable hard drive units.

Media drive transfer units (SCSI or “Small Computer Systems Interface” target devices) have restrictions that limit the overall throughput of the system. This limitation is overcome by breaking the SCSI data stream into a number of separate streams on a frame by frame basis. Each stream is directed to an individual hard drive unit and contains spectral data, frame identification information and ancillary data such as time references, GPS data etc. The downstream software recombines these multiple data streams into a re-collated and ordered flight line.

4. Preprocessing of the *Casi* Data Stream

Data processing includes a number of steps from data acquisition to display. The data, after initial acquisition is backed up onto alternative media. During this process, image frame integrity is checked. The integrity check does not include image quality determination. Following the backup, the image data is converted to radiance data and stored in PCI database format as 16 bit-unsigned integers in spectral radiance units (SRU). Each pixel consists of *n* values, with each value representing the radiance for one of the *n* spectral bands. The principle of the instrument and its operation are described in (Babey *et. al*., 1999).

Conversion to SRUs from raw DN involves a number of steps to correct for instrument characteristics. These steps involve removing dark current generated during the frame exposure, and removal of apparent radiance contributed by the movement of the charge through the sensor during readout (frame shift smear). Each pixel of the sensor has a specific response to light. This is measured and applied in the last step of the conversion process. This entire process is described in (Babey and Soffer, 1992).

The target detection algorithms accept candidate radiance data and compare them against reference spectra. Therefore, to use the target detection algorithms, the system needs to generate a reference spectrum for a specific target. This is known as the “learn”, or “sniff” stage of processing. This is taken as a distinct data set with the data radiometrically corrected, averaged and stored for later use.

Candidate data is collected as a separate step. The data is radiometrically corrected data and presented to the target detection algorithm for comparison with the reference spectra created earlier. This process is known as the “identify” or “fetch” stage. Standard image processing software is used to display the results of this operation to the analyst.

The real-time system requires substantial modifications to the radiometric correction facilities. Many of the algorithms used in the “learn” and “identify” functions are used in the radiometric conversion process. Maximizing the processing efficiency is achieved by combining these functions into one process, thus minimizing data movement, and utilizing many partial results to minimize data operations.

The previously existing output database used the PCI database format, which required an *a priori* knowledge of the output image stream size. The information found in the PCI header, included the image size, which would not be known until the data stream closed. Therefore this format wasn't acceptable for the stream processing needs of the real-time application. The PCI data format was thus replaced with a format that closely resembled the original *casi* data stream. Using this format, the existing display functions could be used and the original data set could be augmented with bands of partial results.

5. Real-time System Design

The overall system design involves a number of considerations:

- Scalability. The algorithms currently used will undergo further refinements. These refinements will likely require more processing power. Improvements in the casi will lead to higher frame rates resulting in higher data rates. More processing power will be required to meet the demands of this system. Data quality monitoring will be a base requirement of future systems. Monitoring will insure that the data is sufficient to meet the needs of the processing algorithms.
- Command/Control. Using a multiprocessor system will require a simple command, control and display function. The casi instrument system has been modified to handle control from an external interface. The other processing systems need to utilize a common interface function from a single operator control system. A standardized network interface needs to be used so that alternative operating systems can be used. The results should be displayed in a straightforward fashion and referenced to the input image data.
- Future hardware improvements. The overall system needs to utilize standard software capable of networking the components at higher data rates. The design should be able to utilize improvements in processors and communications channels minimizing the impact on the software system.
- Data fusion. New instruments being developed by Itres and others should be capable of being incorporated into the system design.
- Commercial applications. The system should be able to operate in a commercial environment. This dictates that size, cost and simplicity of implementation are important considerations.
- Readily available development tools. The systems, specifically the operating system should support a full suite of development tools. Further, the ability to expand beyond the standard services supported by the OS will have distinct advantages for further improvements. Specifically, the ability to utilize custom hardware, improved data communications channels and multiprocessor support will have a dramatic impact on performance.
- Ease of implementation. The time and cost constraints require that a path of least resistance be chosen to implement this system. Many of these constraints will likely be eliminated in the future, so acceptable solutions should not make other major parts of the system useless when the constraints are removed.
- Minimize the impact on the instrument. Modifications to the hardware or software of the casi instrument are considered dramatic tasks. The instrument cannot currently be modified easily to handle new hardware systems due to the use of older software technology. Therefore, the external real-time system should make use of the existing SCSI interface and protocols necessary to communicate with the instrument. Future instrument developments will not have these restrictions.

6. Inter-Process Communications Paths

The real-time system consists of many processes running simultaneously on various operating systems. They all must operate in lock step with the limitations of the instrument in mind at all times. No buffering exists within this system. Hence, the instrument cannot handle data bottlenecks. When this occurs it discards image frames until the bottleneck is cleared. Therefore, the slowest link in the chain will dictate the maximum data flow. The data flow through the system must consider the types of transfer channels that are available.

The data channel types are not limited to hardware channels such as SCSI but also software channels that permit data movement between processes on the same processor. The real-time system consists of a number of distinct processor systems but also some that incorporate multiple processors within a single computational platform. Therefore intra-computer communication is a consideration as well as inter-computer.

With the diversity of hardware and software requirements, a number of different communications paths are used. They include SCSI, Ethernet, communications channels, such as pipes, and connection channels, known as sockets.

7. Implementation Introduction

The implementation system uses two operating systems. A DOS program collects the frame information on the instrument (*casiisvs*) then passes that information via the SCSI interface to another DOS system. This system collects the SCSI information and splits the stream (*separate*) on a frame by frame basis passing the streams to one or more Linux processors optionally replicating the original data stream for backup purposes. The Linux systems perform the radiometric corrections and target identifications (*rtc*) passing their results onto a final Linux system that collates (*collate*) the multiple data streams into a single new stream, then replicates and displays the results.

The DOS System

We elected to utilize a DOS base target mode application to collect data via SCSI from the instrument. This eliminated the need to change the instrument software to use a new interface such as ethernet TCP/IP. (Using a NETBIOS/Samba interface on the instrument was ruled out; the ethernet driver would interfere with the correct operation of the instrument software.) Instead, it was necessary to write a custom target mode application for the data separation function under DOS.

A new program, *separate,* was written whose purpose was to accept the SCSI input frame data stream, then transfer the frame to multiple Linux processors (as many as needed) through the NETBIOS/Samba network over a 100 Mbit/sec ethernet interface. It performs no other data processing functions but is a simple data transfer point. It is able to accept data from the SCSI interface and disseminate the frame data to a number of output devices thus breaking up the original casi data stream into smaller streams permitting a form of coarse parallelism. This process is able to maintain up to 95% of the current casi data throughput.

The DOS to Linux Samba network interface had many problems. These were eventually resolved and the named pipe aspect of Linux was used to temporarily store and hold frame data from the network interface, readying the data for transfer to the LINUX internal processes.

Linux Systems

The Itres utility, *radcorr*, is normally used to radiometrically correct data collected from the instrument in a post collection environment. The raw spectrograph data is typically stored on tape on a custom format hard drive where radiometric correction is later performed as the first step in the target recognition process.

A new program, *rtc* was written as a modified *radcorr* with substantial changes made to support the new network interface, real-time processing, new output data formats and target algorithms. As such the principle functions include:

- The additional ability to read/write from/to a socket device
- The additional ability to read/write from/to a pipe device
- The ability to radiometrically correct raw image data
- The ability to extract ancillary data from the data stream
- The ability to write data in a modified casi stream format rather than the PCI data format
- Inclusion of a number of target detection algorithms to identify targets
- Inclusion of functions that generate target reference spectra for use in detection routines
- Simplification of internal function that improved overall efficiencies in processing

Reading from a simple data file, a single Linux processor operating a Pentium I at 233 MHz (the same speed as the DOS processor) is able to accept a simulated casi data file in DOS and process the data at 95% rated casi speed via the ethernet interface. This did not include target detection functions. However, easy expansion to multiple Linux processes is available if necessary.

Overall Implementation

The following diagram depicts the program model used. It shows the programs, the operating system and the data path used. In this diagram, the logical data flow is shown from process to process. The data physical path, that is the mechanism that the data is transferred, is shown as those arrows connecting with the "switch" device. (The switch device acts as a network traffic controller. The network packets are directed to the segment of the network that contains the destination rather than to all of the segments. Therefore, this reduces the likelihood of a traffic collision between simultaneous transmitters unless the destination is the same in both cases.) The storage utility, *collate*, typically utilizes a standard SCSI interface placing the data onto a local device found on the Linux system. Each processor system is identified as S1, S2, S3 and S4.

In this model, the system is easily scalable by running multiple *rtc* processes, each one obtaining the data from the *separate* process. As the DOS process *separate* is capable of handling 95% of the casi throughput, additional *rtc* processes will linearly scale overall throughput. The Linux process, *collate* accepts the data from all of the *rtc* processes and recombines the frames into a single data stream. The ability to replicate the collated data stream permits simultaneous display and storage.

All of the processors in this system are 233MHz Pentium I units. The data paths are full-duplex 100 Mbit/second ethernet using a data switch to improve traffic flow. Given that *separate* transfers frames to alternate processors, the overall traffic should not collide. However, the same interface that receives data from *separate*, generates and transmits data from *rtc* to the *collate* process. This may generate some collisions between data from *separate* to *rtc* and data from *rtc* to *collate*. This problem did not seem to impact throughput, likely because the data received from *separate* is time-shifted with respect to transmission from *rtc*. To further reduce collisions, we used a full-duplex form of ethernet that separates the physical transmission channel from the reception channel.

The X-Display processor is nothing more that a display system operating the process, *Quiklook.* It displays selected bands be it corrected spectral bands, partial data results or the final target detection band.

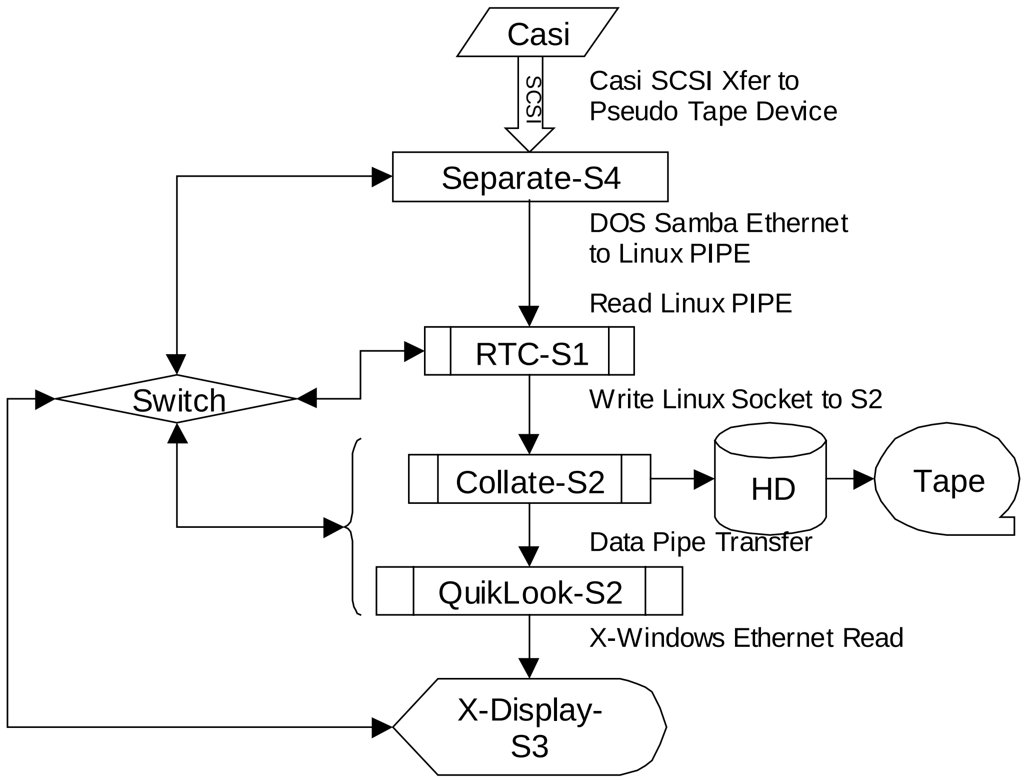

8. Processing Components

Radiometric Correction and Target Identification (*RTC*):

As can been seen in the diagram, the programs *separate*, *collate* and *Quiklook* are really only housekeeping routines. While complex in their own right and a significant effort, they don't perform the actual task of detection. *Rtc* performs the principle functions of data correction and analysis. The two modes of operation learn ("sniff") and identify ("fetch"), are the same program, *rtc*. *Rtc* needs to perform most of the same basic functions in both modes. These include data acquisition, data correction and data averaging. Learn mode additionally transfers a respective sample of the target to a reference file. Identify mode reads, radiometrically corrects and applies the target comparison algorithms to the acquired spectra, using the spectral reference information and generates a new data stream that consists of bands of radiometrically corrected information, bands of partial results and a final band that represents the final result image. Each of these bands is optional and by limiting the output bands, it is possible to maximize throughput.

Learn Mode:

In order to detect the targets, it is first necessary to learn or "sniff" a representative sample of the target. *Rtc*, when run in this mode of operation, sniffs out the sample target averaging the radiometrically corrected spectra across given columns and frames storing the results into a reference file. The reference file can be comprised of multiple target sets. Each set may refer to different targets or environmental conditions such as illumination conditions and target weathering.

Identify Mode:

Rtc's identify mode is the function that analyses candidate targets, searching for targets of interest. In this mode, the system collects and averages dark data, radiometrically corrects the raw image data, and then performs a spectral comparison against all of the data sets found in the reference file using the various target classification algorithms.

A data stream is generated that contains the radiometrically correct spectral data augmented with results of the spectral comparison. The results may contain a number of partial result bands, summary bands etc. The output stream includes:

1. Processing header information
2. Frame ancillary and identification information
3. Radiometrically processed image bands as per standard processing of casi data,
4. Partial result band that represents the partial results for each spectral reference entry. One partial result band per spectral reference set per algorithm is generated in the output stream.
5. Sum of results band. A single channel is generated for all reference entries. This channel is the sum of all of the product band channels normalized by the number of reference entries. Therefore, this one channel can be used to represent the matches for all spectral reference sets. The system can be programmed to only output this one band channel dramatically reducing the total data stream size.

The system can also be programmed so a limited set of bands is produced in the output stream for further downstream processing. This reduces overall bandwidth and processor requirements. The primary utility of generating individual channels is for diagnostic functions.

Data separation system (*separate*):

Data separation is the first component in the data processing path. This is the equivalent to the *casi*'s ability to separate the data frames amongst the various output devices. The *separate* purpose is twofold; it must accept data from the SCSI data stream and transfer the data to named files and it must break up the single data stream into multiple streams. These named files can be either on local DOS drives or networked Linux drives. The networked files can be defined either as normal data files or pipes. As such, this is a simple system of distributing data from the *casi* instrument to the actual high-speed Linux processes. Breaking the data stream into multiple files permits multiple Linux processes to work on the frame data in parallel. This action provided a simple scalable scheme to meet more complex processing needs in the future. If the complexity or the data volume generated by the *casi* became greater than the ability of *separate* to distribute, multiple *separate* processes could be initiated on another processor system.

Data Collation, Replications and Storage (*collate*):

The *collate* process accepts multiple input streams in *casi* frame format and using the information found in the frame header (time), it collates the data back into a single stream generating a composite output stream. This program represents the antithesis of the *separate* process.

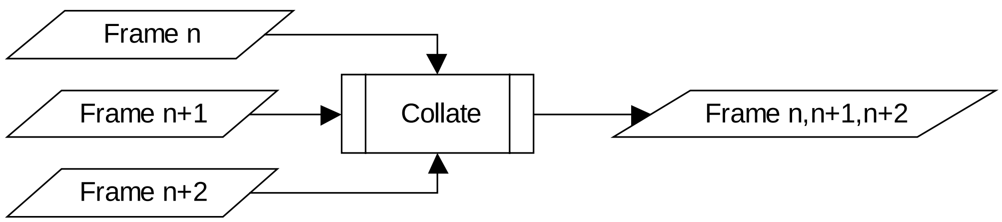

Simple 3 to 1 Collation Process

Multiple *collate* processes may be needed in very complex data streams. *Collate* not only can restore data streams but can also replicate streams. Thus, it is possible for *collate* to transmit copies of the same frame to a number of different processes so that they can each generate partial results and pass these results to a final process combining their results.

Display System (*QuikLook* II):

It was necessary to create a *Quiklook* for the Linux environment and upgrade the program to handle information from pipe data sources using a 16 bit data type. Standard raw casi data is presented as 16-bit integers with a maximum value of 4095. Processed data is presented as 16-bit unsigned integers with a maximum value of 65535.

9. The Field System

The field instrumentation consists of the following:

- Casi II Instrument with no hardware modifications. The software is programmed to utilize a different SCSI target address and limits buffer transfers to 60Kbyte (due to the target DOS memory limitation). This has a minor impact in overall SCSI throughput. The Casi II transfers data from the sensor head through a fully digital interface. An extended length cable was provided to allow for remote mounting of the sensor head from the rest of the support system.
- Rack mounted multiple processor system. This rack contains four complete systems, referred to as **System1**, **System2, System3** and **System4**, utilizing both PCI and ISA bus cards. Each system consists of the following hardware:

    1. 233Mhz Pentium Process card with 10Mbit/second Ethernet, 64Mbytes of memory, video, keyboard, IDE, and SCSI interface
    2. 4Gbyte HD
    3. Floppy Drive
    4. Target mode capable SCSI card
    5. 10/100 Mbit Ethernet

- 10/100 Ethernet Switch

- Removable Hard Drive
- Exabyte 8900 Tape Drive
- Miscellaneous connector cabling and display/keyboard control

The Rackmount systems were used to collect and process the data. **System1**, **System2** and **System3** operated under Linux, and **System4** operated under DOS.

10. Preliminary Field Trials

The real-time processing system was demonstrated for the first time against real mine targets in March 2000. More than 20 different types of real and replica mines were laid out in long rows with 1 m spacing on short and tall prairie grass and asphalt in the Mine Pen enclosure at the Defence Research Establishment Suffield. Weather conditions were sunny with occasional clouds. The *casi*-2 sensor head was mounted on the front of the Threat Detection Group's Mobile Laboratory (figure 1), looking forward and downward at an angle of roughly 45^o. The remainder of the instrument, including all the processing hardware, was mounted inside the laboratory. Reference radiance spectra were obtained by passing different mine types in front of the sensor head, with the instrument in "sniff" mode. The laboratory was then driven slowly (~1-2 km/hr) over the rows of mines, with the instrument in "fetch" mode, and at right angles to the rows, to allow different sun angles with respect to the look direction. Mines were detected in real-time by observing the bright images of the detected mines on the processor monitor. The system was fast enough to compare acquired spectra to multiple (up to 10) reference spectra at the same time (figures 2 and 3).

11. Conclusion and Future Directions

A hyperspectral imager which can detect surface-laid mines in real-time from a land vehicle platform has been demonstrated. To the authors' knowledge, this system represents the first time a hyperspectral imager has detected mines in real-time. The system operated effectively in the field in the land vehicle trials. The real-time system developed to date reflects the time and budget constraints imposed on the prototype design. Most importantly, a speed improvement of two orders of magnitude is still needed for airborne operations. The system was intended for proof-of-concept and as such, it did not have a user-friendly command and control interface. The incident light sensor (ILS) must be improved so that it can be used to convert radiance spectra in real-time to reflectance spectra. (The latter enable classification that is much less sensitive to varying illumination than for radiance spectra.) In addition, a number of other improvements are desirable. These issues are currently in the early stages of being resolved and implemented. The new real-time system unit will consist of:

- Dual high speed Pentium III processors (933 MHz) per bus for a total of four processors per Real Time Unit
- A custom backplane with an asymmetric split PCI bus design supporting two separate processor systems; the entire unit can realize four processors
- Improved DRAM designs
- Dual 100 Mbit/second ethernet interfaces per bus
- Multiple improved SCSI interfaces with significant on-board intelligence
- High speed data switch for the four (2x2) ethernet interfaces with a series of external interfaces to support multiple system units

Significant improvements in the software system will eliminate the need for multiple display systems and provide a mechanism for a common command and control system. DOS will be eliminated in favour of a Linux or similar operating system with full support for ethernet data transmission using custom socket protocols. Replacing the SCSI controller with a custom unit that has a significant amount of buffer memory and intelligence will provide a mechanism to operate under Linux using a custom driver.

Each system unit will have the ability to operate autonomously or linked to other real time system units thus realizing the scalability necessary to meet the detection algorithm requirements. Improvements in both the ethernet and SCSI function will increase the bandwidth between system units.

RAID (Redundant Array of Independent Disks) hard drive storage will be provided through the multiple real-time system units. Each unit will consist of a number of internal high-speed SCSI drivers. External interfaces will permit other system units direct access to these drives.

The *casi-2* instrument is being modified to increase overall throughput and flexibility in the readout structure. The bandwidth requirements to operate the real-time system in the air are substantial, placing a significant burden on the real-time processing platform.

Beyond the immediate role of mine detection, the real-time processing system will provide benefits to *casi* users with more traditional remote sensing applications. It will lead to new ways to reduce overall operating costs, increase reliability and allow for new applications, which require real-time results. Reduced operating costs and increased reliability are achieved by:

- Insuring that data quality requirements of any mission are being met by monitoring the image data, ancillary GPS and IMU data and instrument health. This reduces the number of re-fly lines especially after the mission day has been completed.
- Insuring that the flight lines are being piloted within the boundary conditions set by the mission planners eliminating the risk of void space within the imagery
- Executing pre-processing steps while in flight, thus reducing the time spent on the ground leading to quicker turnaround time to the client. Many steps involved in pre-analysis can be performed during data acquisition. These steps can be performed in parallel eliminating the need to perform over-night processing reducing the number of personnel in the field and eliminating fatigue due to overwork by the instrument operators
- Reducing operational mistakes. The system monitors the data quality and can compare the results against predefined mission parameters
- Executing data replication during flight, thus protecting the flight cost investment if a failure should occur

References

Achal, S., McFee, J. and Anger, C., “Scatterable Mine Detection Using a CASI”, *First Int. Airborne Rem. Sens. Conf.,* Strasbourg, France, September, 1994.

Babey, S.K. and Anger, C.D., “A Compact Airborne Spectrographic Imager (casi)”, *Proc IGARSS* ’89, 12th Canadian Symposium on Remote Sensing, Vol. 2, pp. 1028-1031, Vancouver, July 10-14, 1989.

Babey, S.K. and Soffer, R.J., “Radiometric Calibration of a Compact Airborne Imaging Spectrometer”, Can. J. Rem. Sens., Vol 18, No. 4, pp. 233-242, October 1992.

Babey, S.K., Anger, C.D., Achal, S.B., Ivanco, T., Moise, A., Costella, P.R., DeBliek, J., “Development of a Next-Generation Compact Airborne Spectrographic Imager: casi-2”, Fourth International Canadian Symposium on Remote Sensing, pp I-229, Ottawa, June 21-24, 1999.

McFee, J., Aitken, V., Chesney, R., Das, Y. and Russell, K., "A multisensor, vehicle-mounted, teleoperated mine detector with data fusion", *Detection and Remediation Technologies for Mines and Mine-like Targets III,* Proc. SPIE, April, 1998.

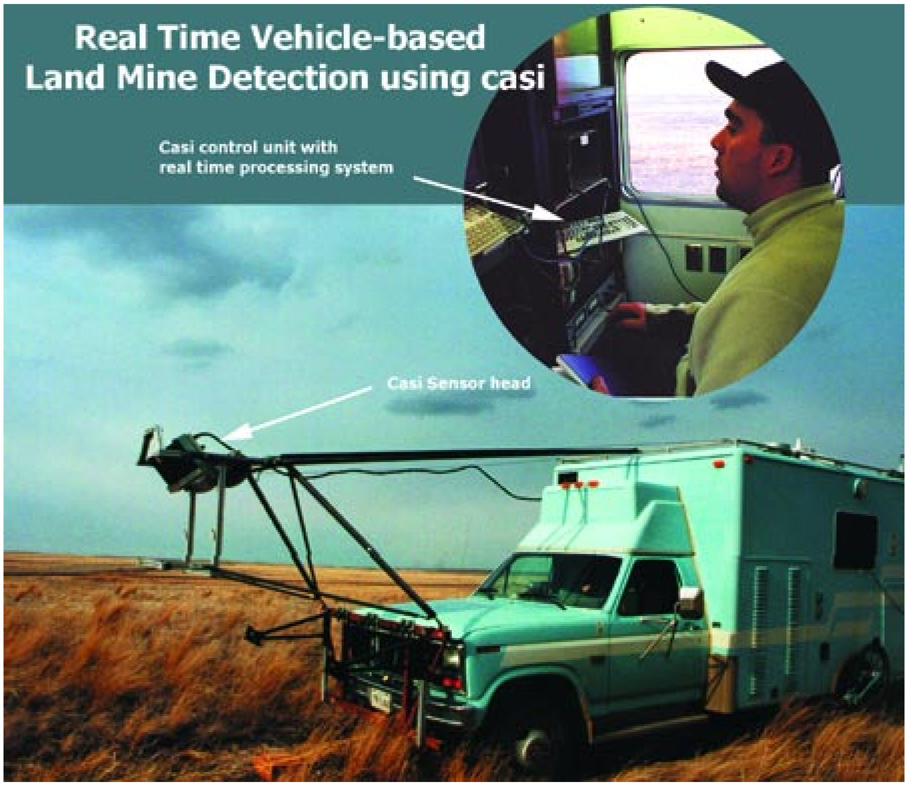

Figure 1: Mobile real-time surface-laid mine detection laboratory

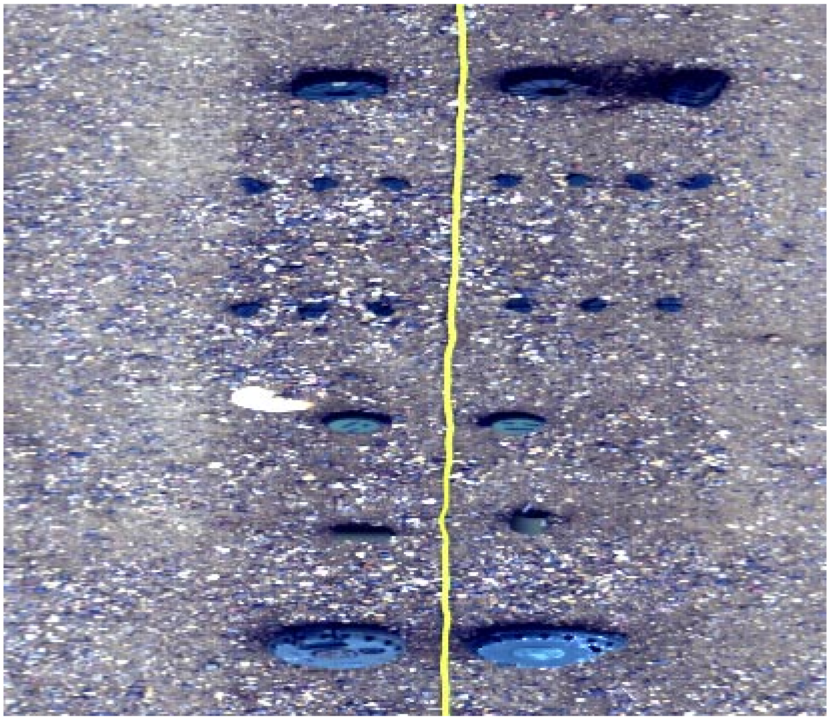

Figure 2: Simulated Mine Field

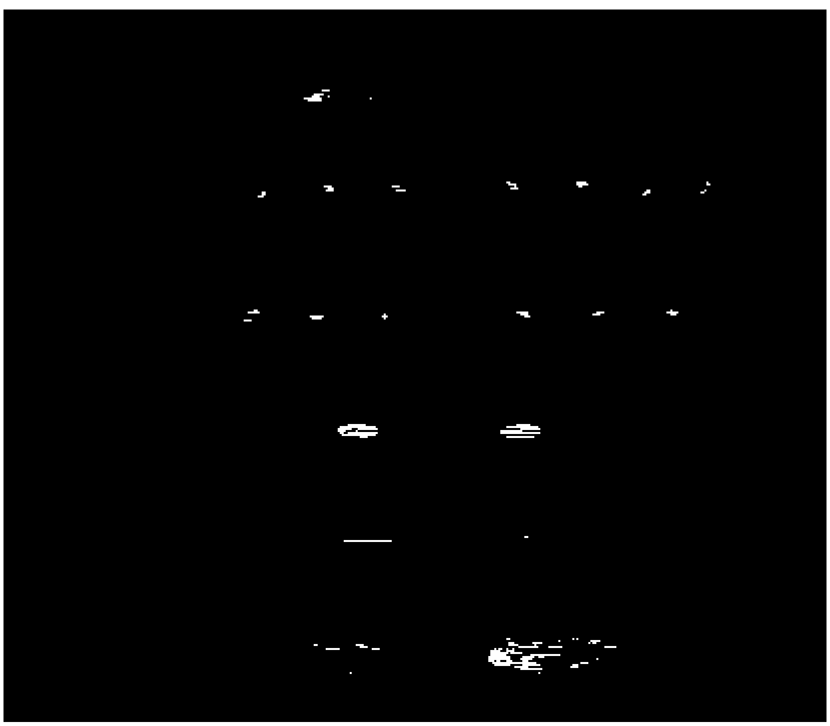

Figure 3: Typical output from the real-time system. Pixels classified as mine-like are white.