Reprinted from

Symposium on Machine Processing of Remotely Sensed Data

June 27 - 29, 1979

The Laboratory for Applications of
Remote Sensing

Purdue University
West Lafayette
Indiana 47907 USA

IEEE Catalog No.
79CH1430-8 MPRSD

IMAGENET - AN IMAGE ANALYSIS NETWORK

P.R. PEARL
DIPIX Systems Limited

I. INTRODUCTION

The launching of the first LANDSAT satellite in 1972 suddenly provided us with a vast amount of relatively cheap remotely sensed data. It was hoped that this freely available data would initiate a dramatic improvement in our ability to map and to monitor the earth's surface. A considerable amount of effort has been expended in demonstrating the potential effectiveness of LANDSAT in solving some of the more pressing requirements of the various resource desciplines. In spite of the apparent success of some of these demonstrations, the operational use of LANDSAT data has not lived up to expectations.

Many factors have contributed to this disappointing performance in particular the long time delay between data acquisition and dissemination, the non-availability of geometrically corrected data and the high cost of effective data analysis and interpretation.

Continuing investment into pre-processing equipment on the part of the responsible government agencies will largely alleviate problems associated with data quality and availability, but unfortunately the cost of image analysis will continue to be born directly by the end user.

The launching of LANDSAT D, SPOT and other remote sensing satellites will further exacerbate this situation, with an increase in both the quality and quantity of remotely sensed data. Consequently, an effective solution to this problem must be found within the near future.

This paper describes an approach which would result in a significant overall reduction in the cost of image analysis by spreading the capital outlay neeeded for a practical image analysis system over a number of end users.

II. IMAGE ANALYSIS SYSTEMS

The current generation of interactive image analysis systems (IAS) may be categorized as follows:

1. Powerful, minicomputer based systems augmented by special purpose high speed hardware. These systems are relatively expensive ($400/K up) and require considerable resources for adequate support and maintenance. Even with technological advances, this type of system will continue to be expensive if adequate image analysis, storage and recording capability is to be provided.

2. Smaller, minicomputer based systems with no (or minimal) special purpose hardware. These systems have been produced by a number of companies and should be somewhat cheaper than those of the first class but provide a much reduced image analysis recording and storage capability.

3. Systems utilizing the resources of a large scale computer. These systems are potentially the cheapest means of acquiring an interactive image analysis facility if a large scale computer is already available although, unless the system can be shared among many simultaneous users, it is doubtful whether the execution of image analysis algorithms would represent efficient utilization of these machines.

An aspiring user of LANDSAT data is faced with the prospect of either purchasing a system from one of these three categories, or of sharing an existing facility with other users. Hardware augmented, single user minicomputer based systems, are frequently outside the budget limitations of many potential users and thus they may opt to acquire a small system without special purpose hardware. Unfortunately, such systems with their relatively low throughput will not be able to satisfy any operational requirements. Few users when faced with the last option are prepared to tie up the storage and processing resources of a large machine with interactive single user image analysis tasks. Both hardware augmented mini-

CH1430-8/79/0000-0286$00.75

computer based systems and large scale computing systems would be more efficiently utilized if their resources could be shared among a number of simultaneous interactive users.

Unfortunately, none of the existing systems of this type provide interactive image analysis on a time shared basis. Any investment required to modify these systems (or design any future system) to support interactive time sharing would be difficult to justify unless a sufficient volume of potentially concurrent users would be likely to benefit.

Concurrent usage would be promoted if interactive image analysis could be pursued from within the immediate vicinity of the users facilities, thus avoiding the cost and inconvenience of having to travel to a central image analysis site. A distributed image analysis network with a centrally located image analysis system (IAS) accessible from a number of remotely located interactive terminals, would be an extremely cost effective approach to LANDSAT data interpretation.

1. The cost of acquiring and operating the IAS would be more easily justified in view of the potentially larger number of concurrent users.

2. The ability to access the IAS from within (or close to) a user's facility will promote the operational use of LANDSAT data.

3. The high cost of on line methodological support could be distributed among a number of concurrent users.

4. A distributed system could be used as the nucleus of a more effective image dissemination system.

III. IMAGENET

A novel solution to these problems is proposed which takes advantage of some recent technological advances, in particular:

- The implementation, by a number of common carriers within Canada and the U.S., of public access packet switching networks.

- The dramatic and continuing reduction in the cost of high speed digital memory.

- The recent advances in the power of microcomputers.

These advances have been incorporated into a concept referred to as IMAGENET (Figure 1). At the core of IMAGENET is a fast, powerful "host image analysis system" (HIAS) based upon a computer controller and augmented by a high speed array processor. Images loaded from computer compatible tape (or perhaps HDDT) would be buffered within the HIAS on disk and all image processing would be performed disk-to-disk via the high speed array processor.

A number of microprocessor controlled "Remote Entry Analyzing Colour Terminals" (REACTs) are connected to the HIAS via a common carrier packet switching network. Commands entered on the REACTs would initiate and control image processing tasks within the HIAS.

Raw or processed image data, suitably tailored for image presentation, would be transmitted to a REACT upon request for interpretation and evaluation. Each REACT incorporates a colour display system capable of performing geometric and radiometric manipulations of the displayed image. Point and area identifiers would be interactively generated on the REACT and transmitted to the HIAS for further reference as required.

For this concept to be technically and economically viable, IMAGENET must meet a number of design goals. These are stated and briefly discussed below:

1. The cost of a REACT should be significantly less than that of an IAS providing equivalent capabilities.

The complexity of the REACT described below is on a par with that associated with some of the more sophisticated computer display subsystems. It is projected that the price of a REACT will be somewhere between \$40K and \$80K depending upon the actual configuration purchased. This compares favourably with the acquisition cost of about \$400K for an IAS providing equivalent capability.

2. The time required to transfer an image from the HIAS to a REACT should not exceed 5 minutes and should be considerably less for most applications.

Any data processing network will introduce a delay between remote user entry and system response. This delay originates from two sources. First, the restricted bandwidth and impulse response of the communication network and second, the processing load on the HIAS incurred through asynchronous multiuser support. Whereas, the second of these restrictions can always be alleviated through the provision of a more powerful HIAS, the first is an inherent limitation of the network.

An acceptable response time is difficult to specify, since ideally the network should be transparent to the user. In reality, whereas the response times for non-image data may be kept negligible, the time required to transmit an image through the network can become appreciable. Its effect may be minimized through the use of any of the following techniques:

A. Image Tailoring

The image size and allocation of bits per pixel is chosen to fit the particular application, hence minimizing the amount of data required to represent an image. Thus, for example, to gain a quick overview of a complete LANDSAT scene, the most significant 2 or 3 bits from three bands may suffice, whereas to provisionally evaluate the results of an image enhancement procedure, a small subarea at full colour resolution would be required. A REACT capable of supporting image tailoring has the further advantage of efficient utilization of its available image memory.

B. Image Buffering

A REACT capable of receiving a second image while displaying the first image undisturbed would also help minimize the effect of network bandwidth. This requirement implies a sophisticated image memory architecture in the REACT.

C. Image Information Packing

This is a technique whereby an image is transmitted through the network with information density **decreasing** with time. This technique could be used, for certain applications, in place of image buffering to enable an operator to monitor the arrival of an image at a REACT. He may then either begin image manipulation prior to the complete image having arrived or even halt data transmission when sufficient image information is presented for his application.

D. Image Data Coding

Image data is encoded prior to transmission to remove redundancy, and upon reception at the REACT, decoded for presentation. Various encoding techniques can be utilized depending upon the application. For example:

1. Each unique pixel is allocated a label and transmitted as such. The Look-Up-Table required to generate a full resolution image is also transmitted to the REACT and used there to reconstruct the image at full 24 bit radiometric resolution.

2. Bit slice run length encoding.

3. Delta-modulation encoding based upon interline or interband differences.

Where applicable, these coding schemes can be implemented resulting in estimated compression factors of at least two. A REACT capable of supporting data coding would further minimize the potential effects of network bandwidth upon operating convenience.

Figure 2 shows the relationship between image size and transmission time through the network. The image size is expressed in terms of "number of bits" with some representative examples identified. The network is assumed to provide a 9600 Baud service with an average rate of 9000 Baud maintained during image transfer. It will be noted that for most "typical" images transfer times are less than 5 minutes and for many applications, would be between one and two minutes.

3. The operating costs of a REACT should be no more than those required to support an IAS of comparable power.

Operating costs may be divided into direct operating costs (maintenance, communication, consumables) and hidden operating costs (software support, methodological support, management). With the exception of communication costs, it is evident that the operating cost of a REACT is less than that for an equivalent IAS. This is particularly true of the hidden operating costs since with these services provided at the central site (HIAS) their costs will be shared by all of the REACTs accessing the network.

Communication costs are more complex and are a function of distance between sites, level of network support and amount of usage. For purposes of comparison, it is assumed that the HIAS is located in OTTAWA and accessed by REACTs across Canada via 9600 Baud lines under the "DATAPAC 3000" service. Then each REACT will incur a "station charge" of $417 per month in addition to a data volume dependent charge. This latter charge is a function of location and is listed in Table 1 for a number of representative examples in terms of a "standard image". A standard image is defined as an image requiring 2×10^6 transmitted bits for generation. This would correspond to any of the "typical" images listed in Table 2. Note that in comparison with the cost of image transmission, the cost of all other data transmission (down line loading, prompts, responses) is negligible.

Table 1. Datapac Rates

Location	Cost per "standard image" transmitted to Ottawa (cents)
Calgary, Alta.	90
Edmonton, Alta.	90
Fredericton, N.B.	98
Halifax, N.S.	63
Montreal, Que.	23
Ottawa, Ont.	15
Quebec, Que.	38
Regina, Sask.	87
St. Johns, Nfld.	81
St. John, N.B.	56
Toronto, Ont.	32
Vancouver, B.C.	94
Victoria, B.C.	139
Winnipeg, Man.	80

Source: "DATAPAC NETWORK RATES" November 1977

Table 2. Standard Image Equivalences

A "standard image" is defined as an image requiring 2×10^6 transmitted bits for transmission. This corresponds to any of the following typical images:

505 lines x 506 pixels x 8 bits	No compression
353 lines x 354 pixels x 16 bits	No compression
353 lines x 354 pixels x 24 bits	Pixel label encoding
353 lines x 354 pixels x 32 bits	2:1 compression
505 lines x 506 pixels x 16 bits	2:1 compression
252 lines x 253 pixels x 32 bits	No compression

Note: Image size is maintained approximately square for comparison. Any ratio of lines to pixels yielding the same total number of image pixels would be equivalent.

If we project that a REACT would be utilized 200 days per year, 5 hours per day at an average rate of 10 standard images per hour, then the yearly communication costs (including station charges), would be $8200 for a REACT located in Toronto, $8800 for one in Quebec City, and $14,000 for one in Calgary, Edmonton or Vancouver. This amount represents a cost which would not be incurred by an equivalent IAS but would be more than offset by the reduced expenditure on maintenance and support required for a REACT.

4. A REACT should be expandable so as to support increased local processing if required in the future.

As the cost of array processors, data storage devices and image recorders decreases with time, it may well be cost effective to perform some of the image analysis tasks locally at the REACT rather than centrally at the HIAS.

IV. A REMOTE ENTRY ANALYZING COLOUR TERMINAL (REACT)

A REACT matching the design requirements outlined above is being built by DIPIX SYSTEMS LIMITED as part of an overall "proof of concept" program to be carried out in conjunction with the Canada Centre for Remote Sensing (CCRS). The program's objective is to demonstrate interactive image analysis performed on the CCRS Image Analysis System from a remotely located REACT using the "DATAPAC" network.

A block diagram of the REACT is shown in Figure 3. A DEC LSI-11 is used to control the REACT. It supports the X25 protocol for the "DATAPAC" link, fields all trackball and keyboard interrupts updating the display characteristics accordingly and routes data between the synchronous serial interface and the digital colour display driver or the operator's terminal as required.

The Digital Colour Display Driver has been designed to incorporate many of the most desirable features required for the interactive manipulation of remote sensing imagery. These features include:

- the ability to perform a continuous ZOOM around any cursor selected point. The ZOOM may be restricted to the X or Y axis only if required, thus facilitating aspect ratio changes.

- a directly addressable video memory of 4×10^6 bits expandable to 8×10^6 bits.

- the number of bits per display pixel dynamically selectable on a line-by-line basis within the range 1 to 32 without underutilization of available video memory.

- 525 line colour monitor used as a trackball controlled movable "window" into the stored image. The size and configuration of the stored image is dynamically selectable and constrained only by the size of video memory.

- extensive arithmetic and look-up-table processing between video memory and digital to analogue converters controllable on a pixel-by-pixel basis.

All functions of the Digital Colour Display Driver are controlled by the LSI-11 microcomputer and exercised from a dedicated control keyboard. The operator uses this keyboard to link the trackball to ZOOM, PAN, graphic functions or cursor movement, to restrict trackball activity to the X or Y axis only, or to select output arithmetic functions or LUTs. Since all these functions are under software control, the repertoire of possible additional functions which could be linked to dedicated keys is quite extensive and could include functions associated with, for example, multiple images, split screen or "image chip" manipulations.

The REACT may be augmented by various mass storage peripherals (e.g. floppy disk, hard disk, tape drive) or the LSI-11 microcomputer upgraded to an LSI-11/23 microcomputer to enhance local processing capability. In fact, a REACT configured with a Tape Drive, a Floppy Disk and a LSI-11/23 microcomputer would form a reasonable basis for a stand-alone image analysis system which could later be used as an interactive terminal for a remotely located HIAS.

REACTs would be linked to the HIAS over the DATAPAC network using an LSI-11 based front-end processor to support the data link protocol. This front-end processor is interfaced to the HIAS via two separate ports; an RS232 asynchronous serial port and a DMA port. It channels all communication from the remotely located operator consoles via the asynchronous serial port so that to software running in the HIAS, these consoles appear identical to local consoles.

Image data is formatted and packetized by the front-end processor via the DMA port for onward transmission to the REACT over the DATAPAC network using the synchronous serial link. Additionally, any other data such as area scans or point identification, is passed back to the HIAS via the DMA interface.

A single front-end processor, such as described here is capable of supporting up to 16 REACTs over the DATAPAC (or equivalent) network.

V. THE FUTURE

Rapid advances are being made in data communication and mass storage techniques. With the groundwork in remote image analysis being established by this "proof of concept" program, it becomes possible to address the overall problem of more widespread use of remotely sensed data. This problem was alluded to above but, in essence, it hinges upon three points:

1. The timely dissemination of remotely sensed data.

2. The integration of other geographic data into the remotely sensed data base.

3. The availability of effective image analysis systems with extensive methodological support.

By the mid 1980s, the present system of disseminating data by mail on CCTs to users saddled, in many cases, with small poorly supported image analysis systems will no longer be able to do justice to the quantity and quality of the available remotely sensed data.

A gradual transition to an on-line distribution and analysis system is envisaged. The hub of the system would be a single "Acquisition and Archiving Centre" (AAC). This centre, under federal jurisdiction, would be responsible for the direct, or indirect, acquisition of the raw data, systematic corrections to the data and archiving of the corrected data. Archiving could be arranged on a hierarchical basis with the most recent (or most popular) data stored in an on-line mass memory system while older data is archived on high density digital tapes. Mass memory systems capable of storing 10^{12} bytes are currently available on the commercial market. This number should be compared with the 10^{11} pixels required to provide full coverage of Canada at the current LANDSAT spatial and spectral resolution.

The AAC would be linked to a small number of strategically located "Regional Analysis Centres" (RAC) using, at first, requested data on HDDT transported by courier services, but as demand rises and as communication costs fall, high speed digital communication channels (satellite or microwave) might become economical. Each RAC, (perhaps under provincial or state jurisdiction) would incorporate an advanced HIAS capable of supporting a large number of REACTs linked to the HIAS by a common carrier data communication network (Figure 4).

End users would access the system from a REACT located within their own facility. They would be able to request data from the AAC for transfer to their RAC ready for interactive remote analysis.

Each RAC could develop into a centre of excellence insofar as methodological support to users within its region is concerned and moreover, could incorporate a regional database of relevant geographic and environmental data. In addition, each RAC would include a high quality image production facility consisting of an image recorder (such as a Laser Beam Recorder) and an appropriate photographic laboratory.

Users would be able to monitor and evaluate the progress of image analysis procedures executed on the HIAS and channel satisfactory results to the image recorder for reproduction.

A configuration based upon this concept might well become economically justifiable within the next decade.

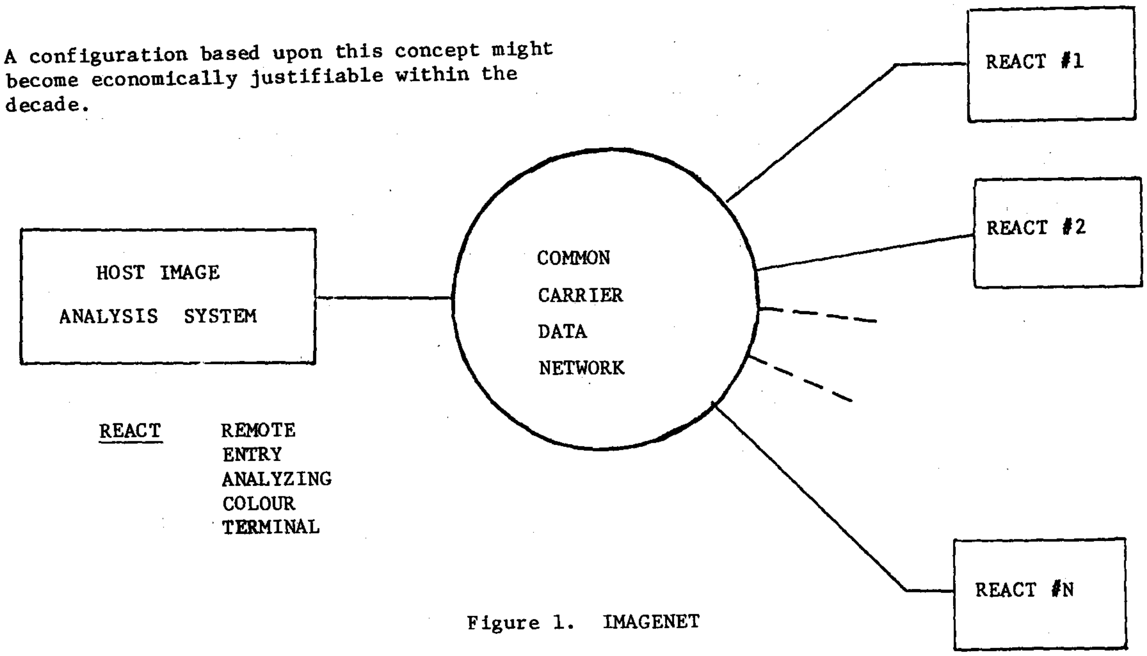

Figure 1. IMAGENET

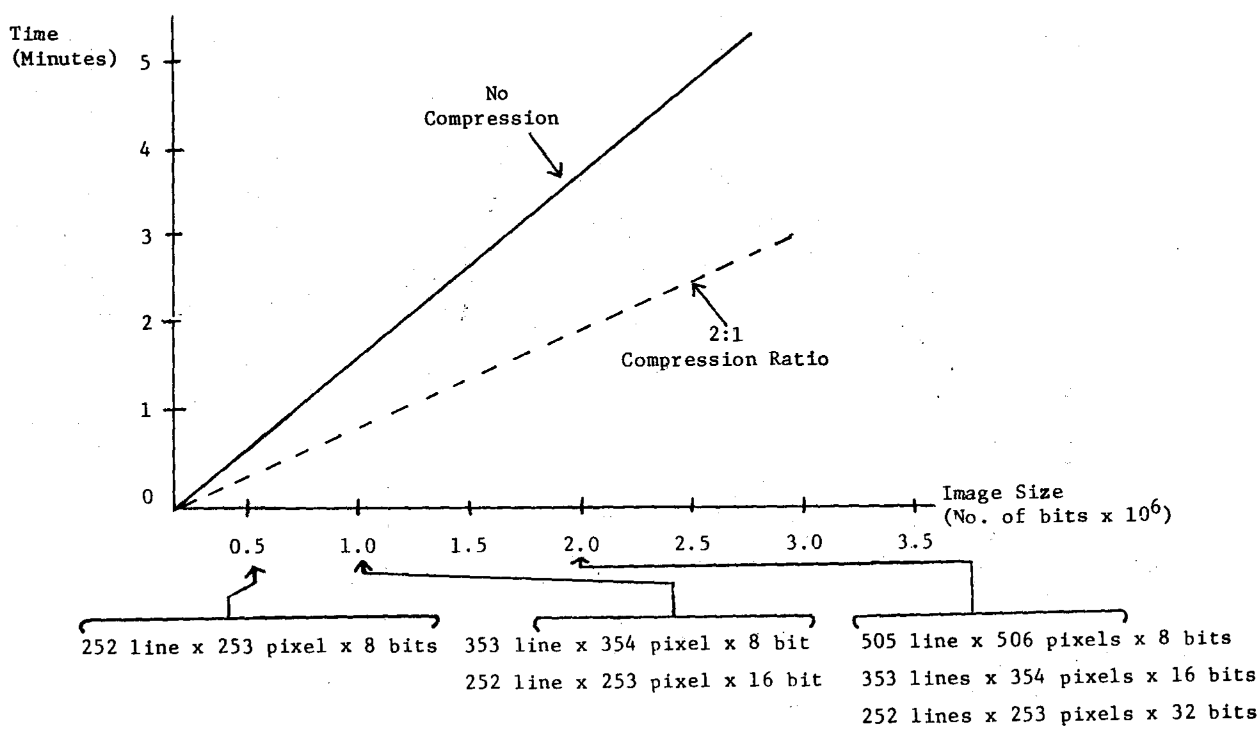

Figure 2. Image Size-Transmission Time Relationship

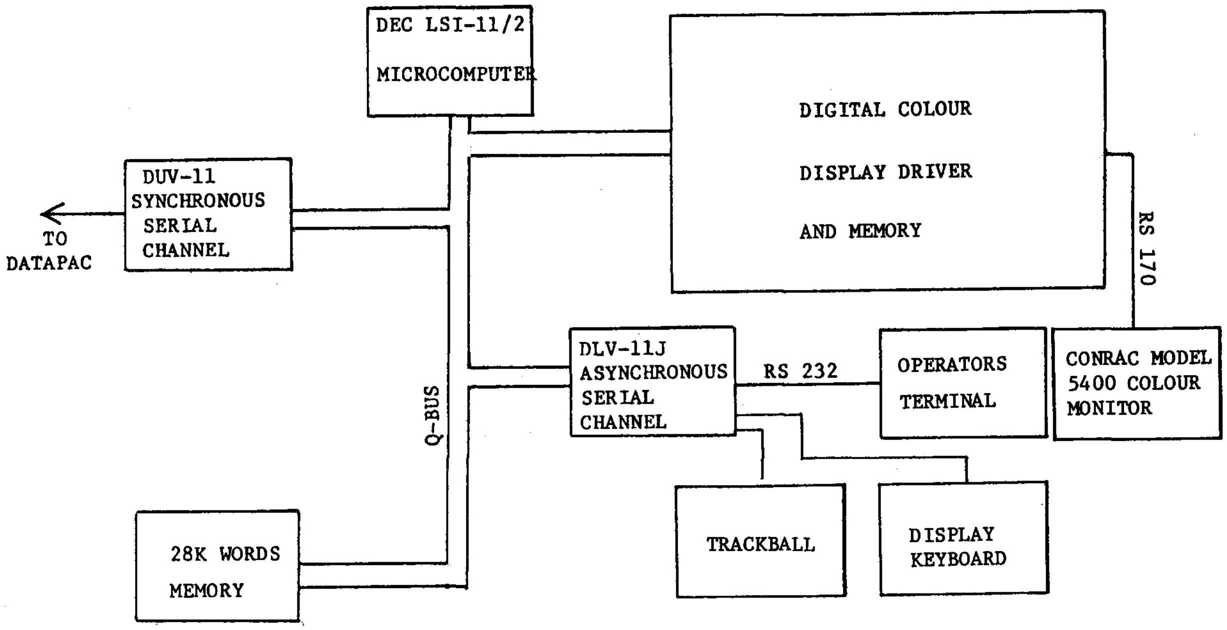

Figure 3. REACT Configuration

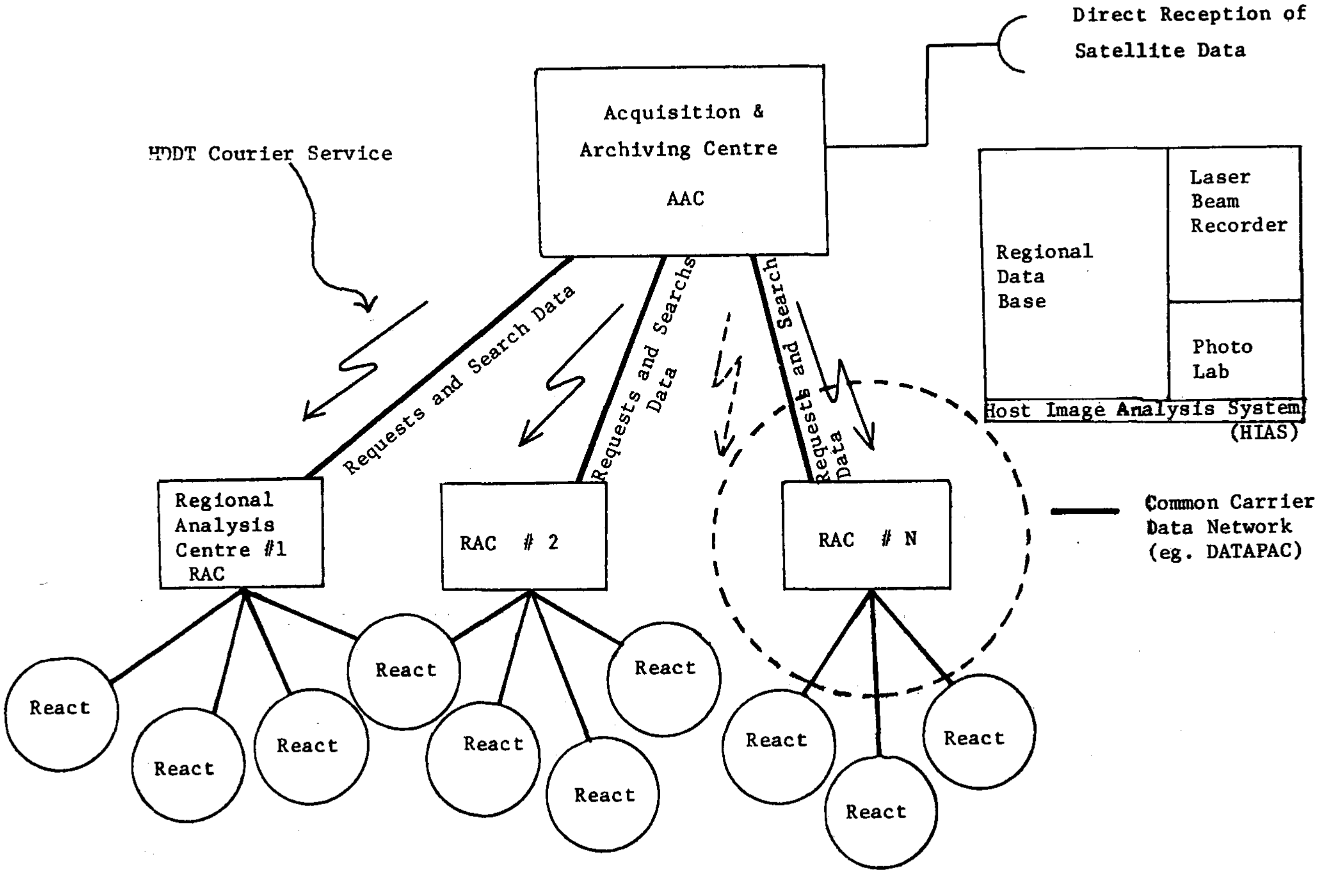

Figure 4. On-Line Distribution and Analysis System

Dr. Paul Pearl is currently president of DIPIX Systems Limited in Ottawa, Canada, a young company specializing in digital image processing systems. He received his B.Sc. and Ph.D in Physics from Monash University in Melbourne, Australia. Following a period of post-doctoral research in plasma physics and lecturing in communications at London University, he joined Computing Devices Company in Ottawa, Canada. There he was a senior engineer responsible for the array processor installation in the Italian Landsat Station at Fucino, Italy and project engineer for the design, build operation and maintenance of the "ARIES" image analysis system.