

PA CLDS
General

AN INTRODUCTION TO THE GEO-INFORMATION SYSTEM OF THE CANADA LAND INVENTORY

by R. F. Tomlinson



CANADA
DEPARTMENT OF FORESTRY
AND RURAL DEVELOPMENT
OTTAWA 1967

AN INTRODUCTION TO THE GEO-INFORMATION SYSTEM OF THE CANADA LAND INVENTORY

by

R.F. Tculinson
Data Processing Coordinator
Canada Land Inventory
ARDA

Published under the authority of
The Honourable Maurice Sauvé, P.C., M.P.,
Minister of Forestry and Rural Development,
Ottawa, 1967

CONTENTS

	Page
Part I. System Requirement.....	1
Part II. System Capabilities.....	2
Part III. System Description Summary.....	8
Data Preparation.....	8
Diagram Showing Flow of Data Preparation Procedures.....	9
Data Input.....	11
Diagram Showing Sequence of File Update Operations.....	13
Data Bank Organization.....	17
Data Retrieval.....	18
Overlay Procedure.....	19
Data Control.....	20
Part IV. Conclusion.....	21
References.....	22
Acknowledgements.....	23

AN INTRODUCTION TO THE GEO-INFORMATION SYSTEM OF THE CANADA LAND INVENTORY

PART I - SYSTEM REQUIREMENTS

The concept of an "information system" is as old as any established discipline. It is essentially a system of collecting and storing information on a particular topic, of examining that information and presenting conclusions based on the examination. The methodology used depends on the type of information that has to be handled and the problems that are presented.

The development of computer-based information systems has become increasingly well known in recent years (1,2,3,4). The idea of a geographic data system to handle information related to a specific location on the earth's surface has been discussed and commented upon (5,6,7,8,9). This paper will discuss the implementation of one such system.

Canada, like many countries, faces an immense problem in both understanding and guiding the development of its land, water, and human resources. One of the major agencies created specifically to implement policy which would attack this problem is the Agricultural and Rural Development Administration (ARDA). A primary task facing this agency is to assemble the data necessary to specifying the problem. The Canada Land Inventory was established to carry out this function.

It was necessary to establish a common basis of data description and to arrange for the physical generation of map and related data, covering: the present use of the land; the capability of the land for agriculture; the capability of the land for forestry; and the capability of the land for recreation and wildlife management. Also, it was necessary to relate these with census data covering socio-economic conditions. Approximately one million square miles is involved in this survey.

If the data are to be gathered at a level where the summaries are directly applicable to provincial and federal resource policy and

regional planning, it is estimated that up to 30,000 map sheets will be generated at various scales.

Lack of trained personnel makes it impossible to examine this amount of data manually in any sensible time, much less to provide a meaningful analysis of the content. A situation is reached where the amount of data precludes their use. The end product of countless hours of survey is often unused. Administrators are not presented with a sound basis for decision making.

A system whereby the map and related data could be stored in a form suitable for processing by a computer, which is also a computer-oriented system capable of rapid measurement and comparison of the data, is clearly desirable. Such a system is the geographic information system of the Canada Land Inventory. The system design and development started in 1963. Implementation began in 1965 and now is in its final stages. Plans call for the system trials to be completed in the summer of 1967, and routine use is scheduled for September 1967.

PART II - SYSTEM CAPABILITIES

The basic capability of the system is that it accepts and stores location-specific information, that is, any information which can be related to the area or place from which it was derived. For example, map information is, by definition, location specific.

Many other types of information, however, are location specific without seeming to be so. Census data, for example, are collected from specific areas of land called enumeration areas, which are recorded on maps. If census data are summarized, the summaries are related to larger, but still specific areas. These are commonly census sub-divisions, townships, counties and so forth. In every case, the data are considered to be homogeneous within the area from which they are collected.

Data can obviously be related to the land surface other than by area. A highway, for example, is a location specific line. A camp site can be thought of as a location-specific point on a map. The geo-

graphic information system will accept all types of location-specific data, and information relating to land resources is most frequently location specific in character.

The geographic information system of the Canada Land Inventory can best be described as comprising two parts. The first of these is the data bank which contains the data. The second is the set of procedures and methods for moving data into the bank, and for carrying out the manipulations, measurements, and comparisons of the data, once there. These two parts will be referred to as the "data bank" and the "information system", respectively. It is quite possible to have the entire geographic information system with full operating capability and have no data in the data bank. The amount of data which can be put into the data bank is infinite, as any number of magnetic tapes can be generated and stored. The data bank can never be "full". Additional data related to any area can be inserted at any time.

The system has the following capabilities:

It will accept maps containing data represented as areas or lines or points. The maps can be of any scale and on any map projection, and they can contain linear distortions. All of these characteristics will be adjusted to a standard format (normalized) when they are put in. Data relating to points only can be put in independent of maps. They are simply related to their latitude and longitude points.

The system compacts and stores information. The compaction is most efficient. For maps at a scale of 1:50,000 with an average density of information it is expected that a complete coverage of the farmed area of Canada (or approximately 600 map sheets) can be recorded on two reels of magnetic tape.

The system can measure any data in the data bank. If the data have been inserted in the form of areas, then each area can be measured. For example, a soil map might be represented by different areas of different soils. In this case, each patch of soil can be measured. If necessary, the total area of any one particular type of soil can be summarized. The area measured of all types of soil can be totalled.

Similarly, the length of any line in the system can be measured. Also, the number of points can be counted.

Of course, measurements over the whole data bank are not performed every time. It is necessary to be able to limit the region from which area measurements are required, or from which line or point measurements are required. This limitation can be effected in a variety of ways. Data can be retrieved within any boundary already described to the system. If, for example, a map of administrative region boundaries has been put into the data bank, subsequent measurements can be carried out within any one of the administrative regions. The lengths of various types of line within the administrative region can be measured. The number of points within any area can be measured. If a desired boundary has not already been described to the system it can, of course, be drawn on a clean sheet and inserted in the normal way. If the desired boundary is not in the system, but is simple enough in shape to be described by a straight line joining points, then it is only necessary to put in the coordinate values of the points.

It will also be possible to limit retrieval by reference to any line already described to this system. The system can be asked, for example, to measure the area of patches of land crossed by the line of a highway. A simple extension of this concept is a band along a line. It will be possible similarly to measure the areas of any specified type of land in a band within two miles of highway, for three miles or four miles or for any desired width of band.

The same capability will exist at and around points. It will be possible to ask the system to provide information at a point or within any specified radius of a point. An example of this might be to provide a measurement of all the areas suitable for sub-divisions within twenty miles of the centre of a city.

A major system capability is comparison of two types of mapped data relating to the same area. Just as two maps can be overlaid to allow the relationships between the data to be examined, the system can overlay any two or more types of data to measure the exact amounts of each type of

land in juxtaposition to the map or maps below.

This will be able to be used as a search capability, whereby a comparison of various types of information is made to find out where a selected set of characteristics occur together on the surface of the land. For example, a request to find suitable landing sites for a helicopter might be made. To locate the site would require an examination of the vegetation map to determine treeless areas, the topographic map to make sure that the area was flat, and the present land-use map to make sure that the area was not populated. These three coverages would be compared to find out where sites with those three characteristics occurred. All points having the desired characteristics would be identified and described. As with the other capabilities, the search can be limited to specific areas or with reference to lines or points.

A further extension of the search capability could result in a "search in context". Referring back to the helicopter landing site example, such sites would be of limited value if, while being perfectly treeless, flat and uninhabited, they occurred as an island in the middle of a swamp. The search routine can be instructed to ignore otherwise desirable sites if they do not occur in a desirable context.

Another search capability that can be implemented is referred to as the "nearest neighbour search". This would be employed when the limit of the search is not definite enough to be specified. The search command would simply request the nearest examples of the desired character to be located. This could be used if a bridge collapsed and it was necessary to know the nearest crossing points on the river. A composite example of some of these capabilities might be an instruction to locate the nearest potash mine which is served by a main highway, north and south railroad connections, and is surrounded by a minimum of 10,000 square miles of good farmland.

Output from the system can be either alphanumeric or graphic. The commonest form of output is perhaps the normal printer output of the regular computer that provides tabular data. In addition to the printer facility will be a graphic plotter which, under the control of the system,

produces a map showing the location of the desired area lines, points, or set of characteristics.

An inherent danger of information systems is the input of data of widely varying reliability, which is assumed to be reliable in subsequent multifactor assessments. The present system can accept a reliability identifier with any type of information. It can be made to keep track of reliability tags so that degrees of reliability would be printed out beside the answer to an information request.

The addition of new data to the system can be done easily and without waiting for large amounts of new data to accrue. Old coverage can be erased and replaced on the magnetic tapes. If desired, both the old and the new coverage can be retained. New survey data at a more detailed scale can be incorporated with previous surveys at smaller scales, provided, of course, that the classifications are compatible.

The advantages of information which is kept up to date, as compared with data which have to accumulate for several years before it is economically desirable to reprint a map, are well known to users of map information.

For many of the day-to-day requests for information encountered by administrators of land resource policy, simple forms exist to allow the administrator to initiate the request without programming assistance. Although the more detailed assessments embodying the full flexibility and capability of the system would best be implemented by someone acquainted with the data formats, a considerable amount of programming effort has been eliminated even at this level by use of pre-programmed modules. It is estimated that with no previous computer knowledge an administrator could be taught to complete normal form-originated requests in one week. Three weeks training and practice thereafter are expected to be necessary for the same administrator to handle more detailed requests. The unusual or very complex requests will need a junior programmer working in conjunction with the system librarian.

In many ways the system is self-monitoring. On accepting a request for information, the first response of the librarian will be to use a system-related "KWIC"* index to check whether that particular request has been made before and, if so, to indicate where the answer is stored in the filing cabinet. If the request has been partially answered, this also is determined. If the request requires new manipulation of data, the system indicates which tapes have the requisite data stored on them.

The tapes then are selected from the library, put on to the computer and the assessment is executed. An extension of this capability is to provide a cost estimate of the work prior to processing. The estimate is based on a preliminary analysis of data density on the requested tapes. Such estimates will be necessary in more complex applications.

The system is independent of peripheral devices such as input scanners or output plotters. While the IBM cartographic scanner is now in use, in conjunction with a D-Mac X-Y digitizer to convert graphic data to digital form, instrumentation is likely to be developed in the next two or three years to combine these functions.

The normalization step, which converts digitized graphic information to the format required by the data bank, is independent of the main system functions and can be changed accordingly.

The system is designed for use on the IBM system 360 Model 50, with 266 K-bytes of storage, 6 tape drives, and 3 disc drives under the control of the operating system. Greater operating efficiency is achieved if the System 360 Model 65 is used. The practical application of the data bank concept and the entire system capability is available by use of this general-purpose computer.

*KWIC - Key Word In Context document indexing and cross-referencing system based on computer sorting of key words in the title. Ref. IBM Pub. E20-8091.

PART III - SYSTEM DESCRIPTION SUMMARY

Data Preparation

Boundary data to be put into the data bank is traced (scribed) on to a clean sheet from the source map (Figure 1). The unique areas or "map elements" are numbered on a transparent overlay and the corresponding classification is transcribed to a data sheet for keypunching.

The traced boundary sheet is placed on the drum scanner (Figure 2), and the scanning operation produces a digitized map of the boundaries on magnetic tape. The drum scanner was developed to meet Canada Land Inventory requirements by the International Business Machines Company. The possible use of the drum scanning approach was first considered in 1963. The preliminary design criteria were established by the Canada Land Inventory in 1964 and development work was contracted to the International Business Machines Company in 1965. The scanner consists of a cylindrical drum on which a map or chart can be mounted and a movable carriage which slowly moves the scanning head across the front of the revolving drum. The scanning system consists of the scanning head proper, its associated electronics and controls leading to a standard IBM 2401 Model 1 magnetic tape drive. The technique employed is to detect the intensity level of light reflected from the map or chart surface and record this information as a series of binary bits written on magnetic tape. The scan head consists of an eight channel device utilizing fibre optics capable of scanning eight scan lines simultaneously. The scanner can accept a map up to 48" x 48" in size. A full size map takes approximately 15 minutes to scan. This includes the time for mounting and dismounting the map. Smaller sheets take a correspondingly shorter time.

It is not within the scope of this paper to give a detailed description of the drum scanner though it is hoped that the engineering aspects will be covered in detail in a future paper. The format of the map-image data on tape is, however, pertinent to the discussion. One map-image record is produced for each .032 inches along the X-axis of a map sheet, and the height of each record area is .004 inches along the Y-axis

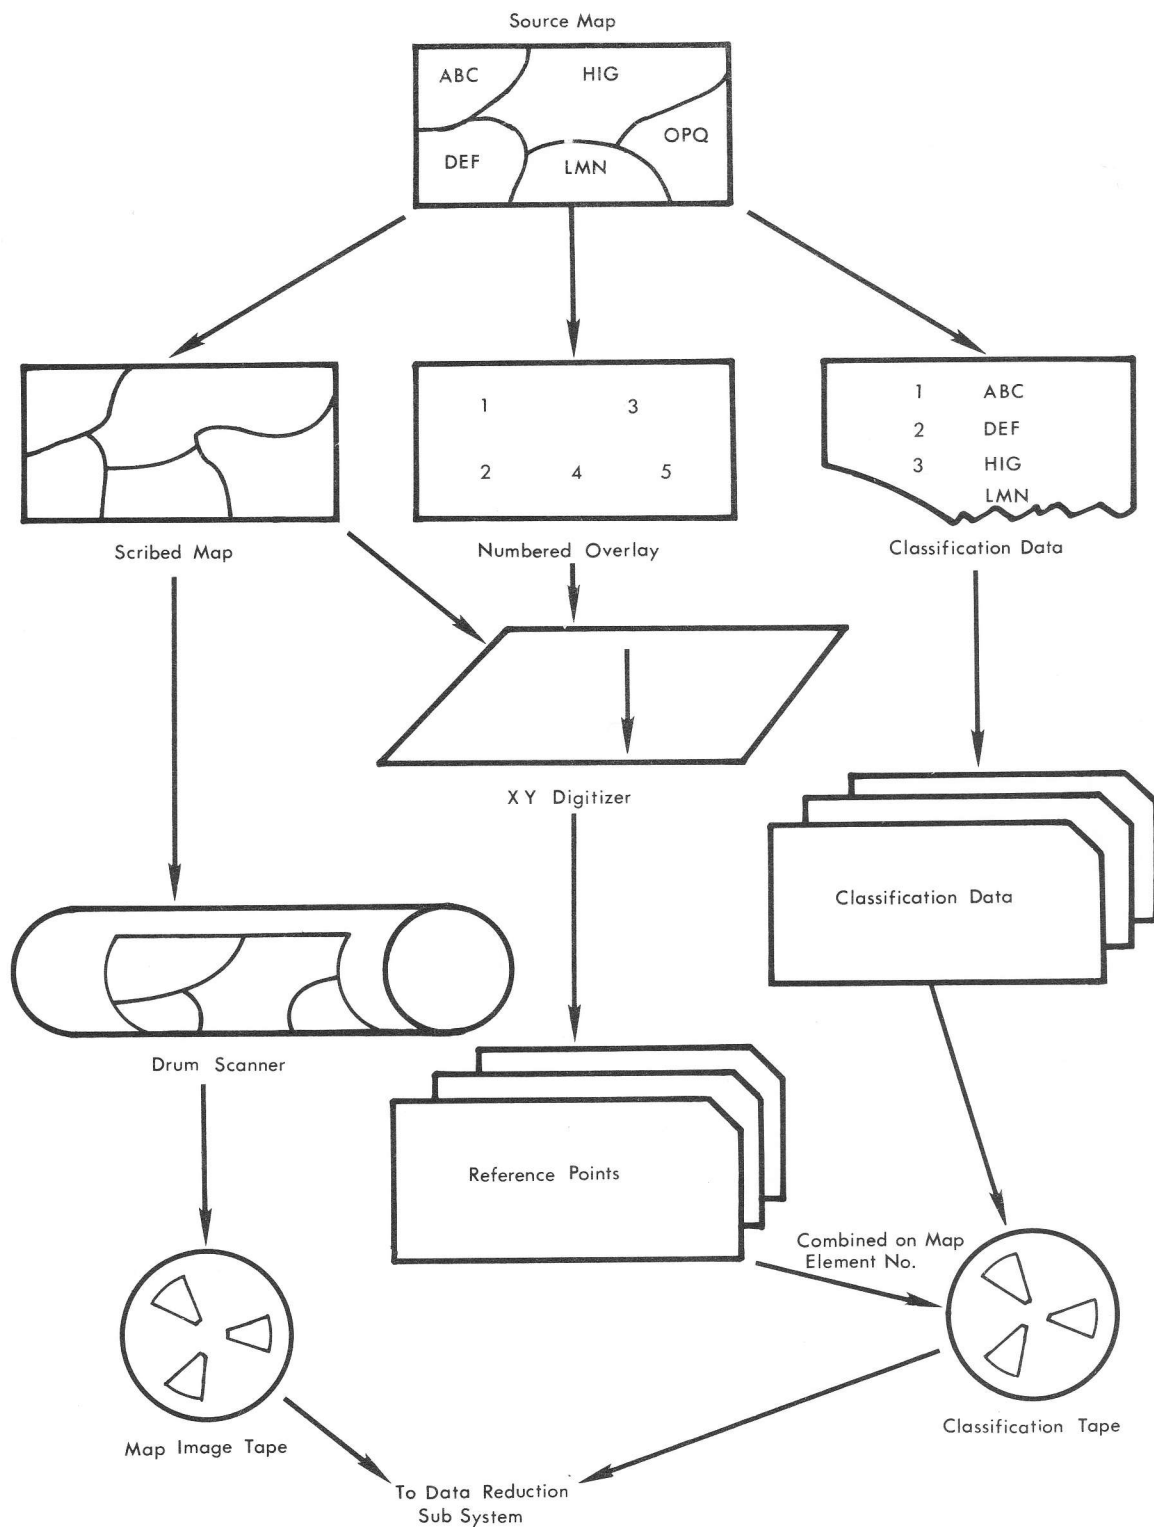


FIG. 1

DIAGRAM SHOWING FLOW OF DATA PREPARATION PROCEDURES




FIG. 2. DRUM SCANNER

The .032 inch record comprising one byte of computer storage is divided into eight bits. Each bit thus represents an area or spot .004 inches wide. Lines usually are scribed .008 inches wide. The scan heads on the scanner are such that if 50 percent of a spot comprises line data, then a '1' bit is generated; otherwise an '0' bit is generated. A line in this manner is represented as a collection of bits which are usually either one, two or three spots in width.

The traced boundary sheet with the transparent numbered overlay is placed on a D-Mac cartographic X-Y digitizer (Figure 3), where the four reference corner points and the coordinates of one reference point per map "face" are digitized. A map face is any one of the distinct areas that together make up the surface of the map. As noted before, information related to a face is considered to be homogeneously distributed within that face. Punched cards represent the output from the X-Y digitizer operation. In the future, the digitizer probably will use an incremental magnetic tape for output if it is found that the error-edit capability of cards is not needed. The classification data sheet is now keypunched, though this also may be taken directly to tape output. Classification data and the digitized reference points are combined on the basis of map face number to result in a classification tape.

Data Input

The basic approach to the input of map data is to reconstruct the points comprising the scanned map-image into line segments, i.e., that part of a line that lies between adjacent vertices, then to combine these segments with the classification information to produce map faces. These are a basic unit of storage.

The following are some of the steps in this input procedure: As a preliminary, the scanner and classification tape volume serial numbers, coverage and map identification, and similar data are put into the subsystem monitor which controls the flow in the subsequent operation. The classification tape is edited for data consistency and is changed into system format during this stage (Figure 4).




FIG. 3. X-Y DIGITIZER

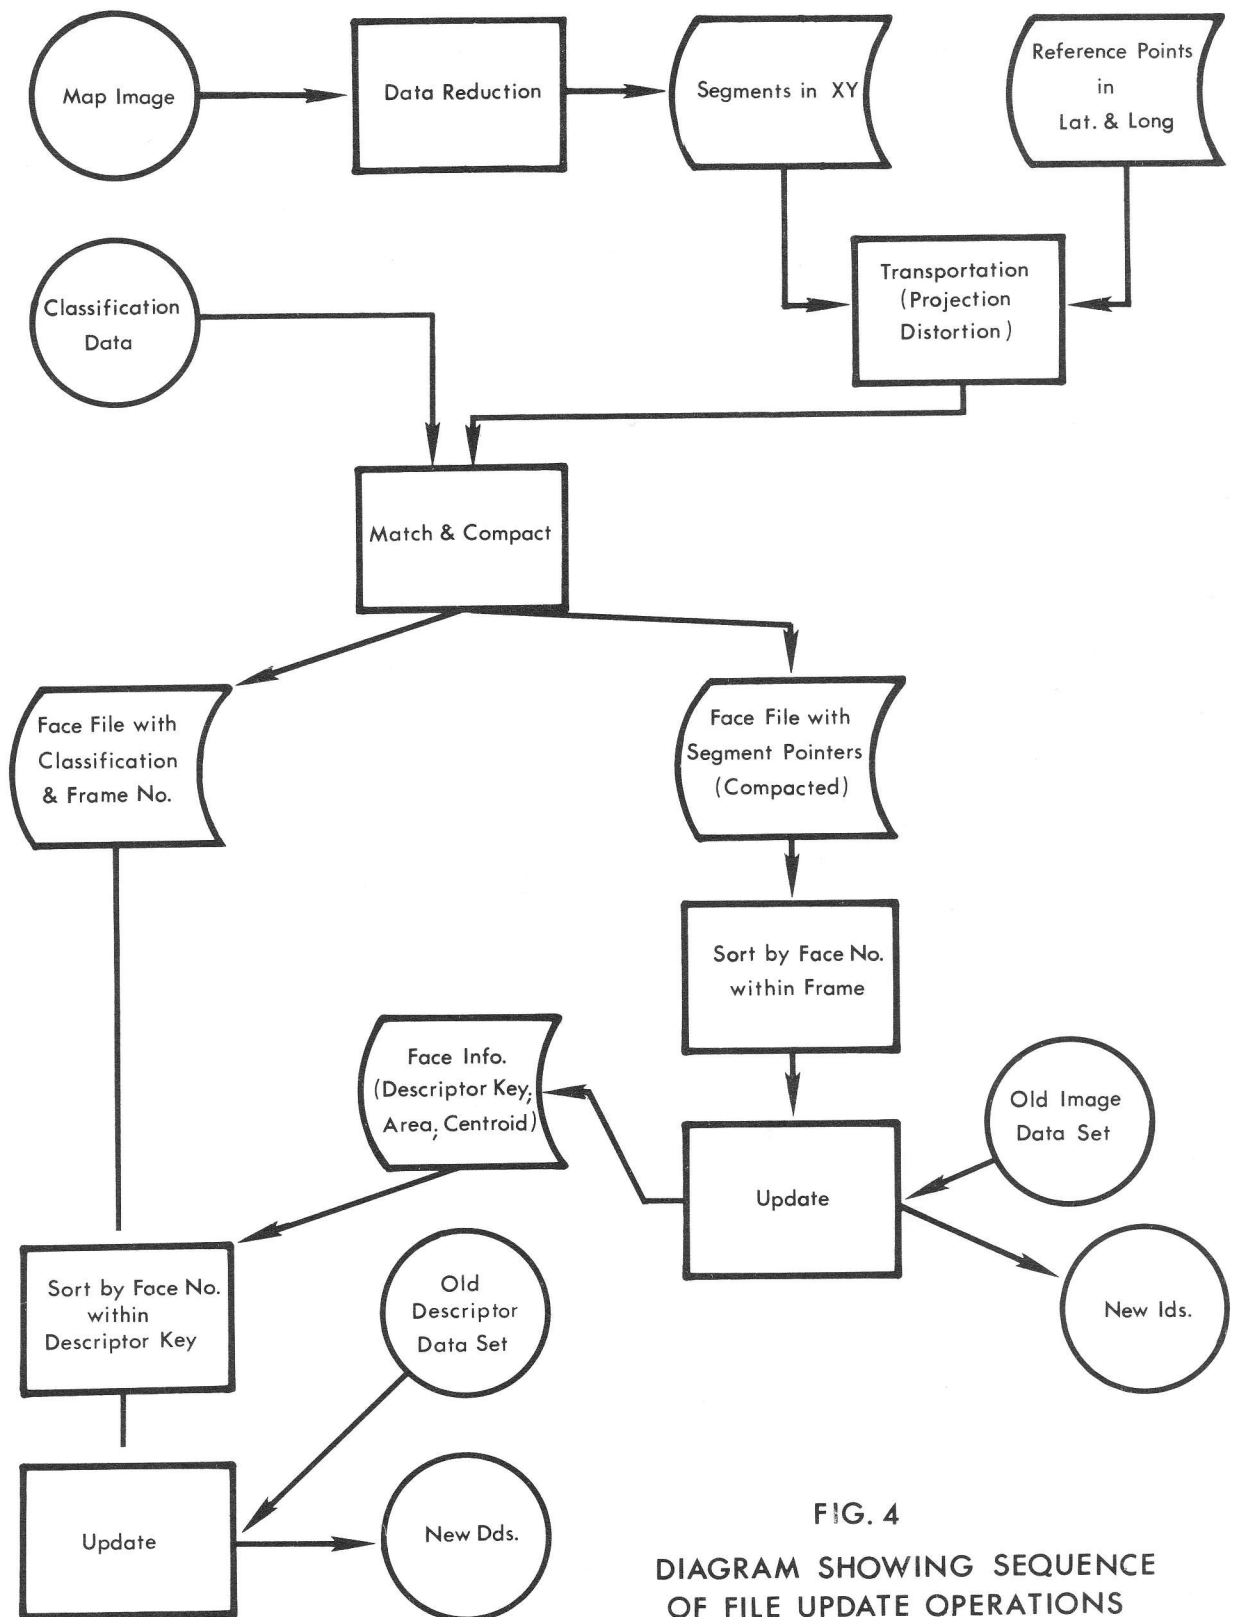


FIG. 4
DIAGRAM SHOWING SEQUENCE
OF FILE UPDATE OPERATIONS

The map-image tape then enters the main map data reduction procedure. Since a 30-inch by 30-inch map generates over 56 million bits, which itself occupies over 7 million bytes of computer storage on an IBM System/360, the data reduction of the map-image is performed sequentially on smaller units known as "sections". The use of a square (or nearly square) section results in considerably longer lines being available from the map for processing at one time than would be the case if a less-than-square rectangle were used. A computer with 256 K-bytes of core storage can handle a section size in the order of one square inch.

Each spot in the cloud of spots which make up the lines is assigned a 'V' value. This is a measure of the number of information-carrying spots surrounding it. This minimizes the effect of stray noise bits and tends to pick out the center points along the line. The search algorithm follows the highest V values; it eliminates the redundant spots in the cloud.

The center points are coded to identify line intersections (or vertices) and the sense of direction of the line. Having thus located the points which comprise boundary lines, it is a simple task to record the X and Y coordinates of each point along a segment.

The system requires that descriptive information be related to map elements. One method of accomplishing this is to apply an identifying tag to both sides of the line. This tag also indicates in which direction the line was first followed, this being necessary if the sides of the line are to have a constant meaning. The identifying tags are called "system colors". They are analogous to the colors in a political map. A sort-and-search of these colors enables segment connecting to be accomplished, and hence faces to be assembled.

Using the reference points in latitude and longitude taken from the four corner points of the map, a transformation is carried out which locates the X-Y digitizer map-element reference points within the scanner. Map projections, which can vary from source document to source document, are normalized. Calculations are made to correct for linear distortion and skewed orientation of the source document on the scanner or digitizer. The

transformed "map-image data set" and the classification (or "descriptor-data set") then are matched and compacted. During this match-and-compact operation, the map-image coordinates are recorded in terms of a standardized geodetic coordinate system. This allows a uniform base for storage and the subsequent measurement and overlay procedures.

The choice of a standard coordinate system was a major consideration. The eventual measurement needs, (i.e. area, length and centroid) required the chosen system to be locally cartesian. However, a coordinate system based on a projection can result in a system of regions, each with its own coordinate system. This problem is particularly pertinent when one considers an area as extensive as Canada.

Careful investigation indicated that a system comprised of the geodetic latitude and longitude had many advantages. The smallest division in the geodetic coordinate system used in the data bank is called a unit grid. It represents an angular displacement of $1/2^{24}$ degrees. This was derived quite empirically. Using a 32-bit word, 8 binary bits allow a span of 128 degrees which is sufficient to encompass Canada. The remaining 24 bits represent the possible subdivision of any one degree.

The theoretical resolution of the system is determined by the actual distance on the ground covered by this unit grid, which at 45 degrees latitude is just over 1/4 inch in the latitudinal (or X) direction. This is considered adequate for the data being put into the system.

Scale within the system is in terms of the unit grid distance. Factors from 0 to 31 have been devised to provide coarser resolution.

To handle map information within the system, it is convenient to subdivide the coordinate system into processable regions called "frames". A frame has an equal angular displacement in the X and Y directions, hence is a square in the geodetic coordinate system. It is not convenient to process frames which contain more than $2^{16} \times 2^{16}$ grid points, therefore arbitrarily the X and Y coordinates are restricted to 16 bits for 1/2 word with respect to the frame in which they fall.

A relatively simple calculation reveals that a map of average density (30 inches by 30 inches, with 800 inches of boundary lines), will occupy 200,000 bytes of storage if no scale change or transformation is performed, and if 1/2 word is taken for each X and Y coordinate. With up to 30,000 maps envisaged as the primary content of the data bank, a compact notation for storage of coordinates was essential.

With a code based on direction change between coordinates and distance between coordinates, a sequence of two-bit codes can be used to describe coordinates. A sample line, requiring 864 bits for normal X-Y recording, occupies 76 bits in compact notation. Thirty-two of these are taken by the X and Y coordinates of the start point. If required, lines with regular patterns can be further compacted by storing the pattern and an indication of how many times the pattern is repeated.

In the match-and-compact phase, routines are carried out to calculate the area of each face, the centroid of face elements, and the length of line elements. In the same phase, an extensive error analysis is performed to ensure that the map is topologically correct. Errors found at this stage are documented by a series of error messages on the computer printer.

The match-and-compact operation produces two index files. The first of these is a face file with classification and frame number which, when sorted, is used in updating the descriptor-data set. The second is a face file with segment pointers which is used to update the image-data set. Incorporated in the second file is the basic compact notation of coordinate data by frame number. The update routine for the image data set provides the geodetic properties (area, centroid and length) as required by the descriptor-data set update. Both of these update routines can produce error listings as new data are matched with data in the data bank. Again, error corrections are carried out as an update to the primary map-data reduction phase.

A systems approach is essential to error correction and this will only be resolved by trial with a working system. Given a high percentage of errors requiring reference back to source documents or even to field

survey, the relatively expensive man-machine interaction which uses cathode ray tube displays would add little, if anything, to the efficiency of the error-correction procedure. On the other hand, given a high percentage of errors of a strictly cartographic nature, and not requiring reference to source documents, the cathode ray tube approach, which enables such errors to be displayed and immediately corrected with a light pen, would have considerable merit. Both approaches will be investigated during the system trials.

Data Bank Organization

The data bank is split into classification data contained within the descriptor-data set and boundary data contained within the image-data set. Three levels of file organization are envisaged. These are:

1. Consecutive;
2. Regional;
3. Indexed.

These file organizations, together with an unstructured or structured version of the classification data within the descriptor-data set, have been combined into six levels. Five of these will be possible within the present scope of the data bank.

Using the descriptor-data set as an example, the relationship between the various levels can be thought of as follows: Level 1 represents the basic descriptor-data set arranged by consecutive face number; Level 2 represents a sorted Level 1, grouped according to some selected characteristic or set of characteristics; Level 3 is the equivalent to Level 1 for a specific region or group of regions; Level 4 can be thought of as a Level 3 which has been structured by grouping the faces relating to a certain characteristic or set of characteristics; Level 6 is a Level 2 or 4 which is not only structured but has an index of its contents available to facilitate further search. Level 5 is not implemented as an indexed consecutive file is not an advantage.

In the descriptor-data set for each map element, there is a list of pointers to the frames containing relevant parts of the boundary

information for that map element. The format of this key varies with the level of file organization, but in all cases, it serves to relate the image-data set to the descriptor-data set. The record formats of the various levels of descriptor-data set are illustrated below.

Level 1, 3

Record Type	Cover- age Number	Map- Element	Geodetic Data	Factor Data	Frame List	Level 3 Region List
----------------	-------------------------	-----------------	------------------	----------------	---------------	---------------------------

Level 2, 4, 6

Record Type	DDS Key Classification Data	Geodetic Data	Non-Key Factor Data	Frame List	Region No. List
----------------	-----------------------------------	------------------	---------------------------	---------------	--------------------

Data Retrieval

With the exception of the actual boundary coordinates, all the descriptive information for a map element is contained in the descriptor-data set. Most data retrieval procedures thus are expected to apply only to the descriptor-data set. File organization of the descriptor-data set is designed to facilitate the formulation of a data request. The data set describes coverage in terms of compiler-language declaration statements.

The required retrieval of assessment procedures are combined with the proper data-specification statements as input to the compiler. The compiled object program is combined with input/output sub-routines to form an executable procedure. This approach eliminates many of the interpretive procedures (and the necessary error-diagnostic routines) that would need programming if the same high degree of data flexibility was desired. A compiler-produced operation can be more efficient than a similar interpretive procedure.

The system also is able to accommodate new types of data or new methods of handling data with great facility. A careful consideration of the available compilers and languages has indicated that PL/1 offers the

greatest capability for data description and data handling. Thus the data specification statements of the descriptor-data set are in PL/1 language. All requests for data retrieval ultimately are entered into the system in PL/1 language. Extensive use is made of macros in modular form to handle commonly recurring low-level requests, particularly those within the descriptor-data set. Higher order requests (particularly those incorporating the image-data set and the overlay function) must be written in PL/1. Even in the higher levels, however, the interface between the user-written procedure and the data bank proper will be facilitated by a system of input/output macros.

Overlay Procedure

The overlay capability of the system is basic to the comparison of any two sets of location-specific information which have different data bases. This is the well-known function of putting one map over another and examining the resultant data relationships. The overlay procedure can be thought of in three stages:

- (1) The normalization of the input data;
- (2) The map element overlay;
- (3) The descriptor overlay.

As with the manual method of overlaying maps, the image-data sets to be overlaid within the system must be at a common scale and frame size. The input selection phase of the overlay procedure accomplishes this objective.

The map-element overlay operation is similar to the primary map-data reduction phase. The line segments must be arranged, the intersections must be identified and the new map elements must be created. This is accomplished by a procedure which uses a 1-bit-per-point frame core image. The base frame is put into core and the coordinates of the intersections of this base frame are put into a vertex table and sorted Y with X. The overlay frame is treated similarly, and the coordinates of the new intersections are noted in the vertex table. If more than two frames are involved, the subsequent frames are treated similarly.

When all frames have been entered into the vertex table, it is merged to one table and the frame information is re-entered to identify the segment coordinates of the combined frame. New segments are used to create new face elements and are recolored. A correlation of colors is carried out for those face elements which span more than one frame. This is accomplished by a sort on the face-element system color. A re-sort to frame sequence establishes the combined image-data set.

To incorporate the descriptor-data file for a Level 1 coverage, the new combined face-element record is exploded to produce a separate record for each coverage included in the overlay. The record containing the map element numbers of these coverages is then used to extract the classification information for the combined face-element. A re-sort to the new face-element sequence creates the new descriptor-data set.

Data Control

Data control within the system is achieved by the system monitor. The system monitor accepts pertinent data on the history of map-data manipulation within the system at all times. Many of the responsibilities for system control in such an open-ended system must rest with the system librarian.

The librarian's responsibilities include deciding whether coverages are permanent or temporary, selecting the resolution at which boundary lines for various coverages need to be stored, and deciding the structure levels of the descriptor-data sets and of the pattern for restructuring. He is also responsible for providing the procedures which edit the classification data in the preliminary phase of the map-data reduction sub-system. He must tailor the data retrieval modules to efficient, specifically applicable retrieval requirements. He is in control of the flow of individual maps within the system and, similarly he must evaluate the practicability of assessment requests, including the avoidance of duplicate assessments.

PART IV - CONCLUSION

The Canada Land Inventory geographic information system has several new concepts and techniques. The over-riding consideration, however, is the creation of a data bank containing not only descriptive information, but a compact, useful form of the related boundary information. The system can be of inestimable value to a country such as Canada in providing a sound basis on which decisions on planning and development of resources can be made particularly in the context of regional rehabilitation and project planning. The system is basic to a geographical understanding of the country and has application in any nation where the developing economy is concerned with the natural resources.

REFERENCES

1. Berkeley, Edmund C. 1962. The Computer Revolution. Doubleday and Company, Garden City, N.Y.
2. Vickery, B.C., 1963. On Retrieval System Theory. Butterworth and Company, London.
3. Becker, Joseph, and Hayes, Robert M. 1963. Information Storage and Retrieval Tools Elements Theories. John Wiley and Sons, New York, N.Y.
4. Henderson, M.M., Moats, J.S., Stevens, M.E., and Newman, S.M., 1966. Co-operation, Convertibility and Compatability Among Information Systems: A Literature Review. U.S. Department of Commerce. Pub. 276.
5. Garrison, W. L., Alexander, R., Bailey, W., Dacey, M.S., and Marble, D.F. Data Systems Requirements for Geographic Research. Tech. Report No. 1. ONR Task 1 No. 389-142; Contract NONR 1288 (35).
6. Dacey, M.F., and Marble, D.F. Some Comments on Certain Technical Aspects of Geographic Information Systems. Tech. Report No. 2. ONR Task 1 No. 389-142; Contract NONR 1288 (35).
7. Coppock, J.T. 1962. Electronic Data Processing in Geographical Research. Professional Geographer 14 (4), 1-4.
8. Coppock, J.T., and Johnson, J.H. 1962. Measurement in Human Geography. Economic Geography. 38 (2), 130-37.
9. Roll, L.L. 1966. Geographic Data Processing. Photogrammetric Engineering, 32 (6).

ACKNOWLEDGEMENTS

The author gratefully acknowledges the support and encouragement given this work by ARDA, particularly from A.T. Davidson, L.E. Pratt, and W.A. Benson, Chief, Canada Land Inventory, and from the International Business Machine Company of the IBM/Ottawa staff associated with the programming and computer applications of the project; specifically: Guy Morton, Frank Jankulak, Don Lever, Art Benjamin, Robert Kemeny, Peter Kingston and Bruce Ferrier.