TR
693
.L83
1985

AUTOMATED DATA ORDERING IN PHOTOGRAMMETRY

By
JAMES R. LUCAS
National Charting Research and Development Laboratory

PREPARED FOR PUBLICATION IN A
SPECIAL COMMEMORATIVE VOLUME
BEING ISSUED BY THE TECHNICAL
UNIVERSITY IN ZURICH, SWITZERLAND,
HONORING DR. HELLMUT H. SCHMID

U.S. DEPARTMENT OF COMMERCE
National Oceanic and Atmospheric Administration
National Ocean Service
Office of Charting and Geodetic Services
Rockville, Maryland

MAY 1985

TR
693
.L82
1985

AUTOMATED DATA ORDERING IN PHOTOGRAMMETRY

By
JAMES R. LUCAS
National Charting Research and Development Laboratory

PREPARED FOR PUBLICATION IN A
SPECIAL COMMEMORATIVE VOLUME
BEING ISSUED BY THE TECHNICAL
UNIVERSITY IN ZURICH, SWITZERLAND,
HONORING DR. HELLMUT H. SCHMID

U.S. DEPARTMENT OF COMMERCE
National Oceanic and Atmospheric Administration
National Ocean Service
Office of Charting and Geodetic Services
Rockville, Maryland

MAY 1985

AUTOMATED DATA ORDERING IN PHOTOGRAMMETRY

James R. Lucas
National Charting Research and Development Laboratory
Charting & Geodetic Services, National Ocean Service
National Oceanic and Atmospheric Administration
Rockville, Maryland 20852

ABSTRACT

When the data in a photogrammetric block-bundle adjustment is ordered in the most advantageous manner the resulting normal equations have a banded structure for which both storage and computer time is near minimal. Several algorithms designed to reduce bandwidth, profile, or fill are compared using a hypothetical photo block in search of an automated method that would improve on the manual process of cross-strip numbering. None of the algorithms investigated could be considered superior to manual sorting in all respects, and therefore, a heuristically derived algorithm is proposed that will duplicate optimum manual sorting in an ideal case.

INTRODUCTION

When Dr. Hellmut Schmid went to the Technical University in Zurich, Switzerland, as a Visiting Professor in January 1974, it became necessary to hand down some of his many active projects to various members of the Geodetic Research and Development Laboratory (GRDL) of the National Geodetic Survey (NGS). One such project, the simultaneous adjustment of more than 1,200 metric quality photographs obtained on the last three Apollo missions, for the purpose of establishing a Selenocentric control network (Doyle et al, 1977), was temporarily assigned to the author. In September 1974, when Dr. Schmid retired from the NGS and moved permanently to Switzerland, this temporary assignment continued to be a challenging and rewarding responsibility.

Selection, identification, and measurement of the imagery were accomplished by the Defense Mapping Agency, which also provided the orientation angles of each photo, as determined from the coupled stellar exposures. The MUSAT IV program (Elassal et al, 1970) was provided by the U.S. Geological Survey, which was cooperating in the project, so there was very little software development required. The only major obstacle that remained was structuring the data such that the adjustment of more than 23,000 unknown parameters could be fit into the available computer, a CDC-6600 with approximately 300K words of available storage, and completed within a reasonable time.

Cross-strip numbering of photographs to minimize the bandwidth of the normal equation structure (Gyer, 1967) was common practice in dealing with conventional photogrammetric networks.

However, the varying amounts of overlap that resulted from orbital photography within each mission and the criss-crossing of strips from missions with different orbital inclinations provided a more difficult problem than had been anticipated. Attempts to order the photos for acceptable bandwidth consisted of visual inspection of a ground track graphic by experienced photogrammetrists and a more elaborate scheme based on sliding a template perpendicular to a line approximating the long dimension of the block and selecting photos in the order in which their plotted nadir points were encountered. The best effort resulted in a bandwidth of 504 unknowns (84 photos) which would have required more than 250K words of storage for that portion of the normal equation matrix that must reside in core.

The ordering that was finally used to accomplish this adjustment was provided by the U.S. Naval Ship Research and Development Center's BANDIT Program, which employed the algorithm of Cuthill and McKee (1969) and had been implemented on the NGS computer by Robert H. Hanson of the GRDL. The resulting bandwidth of 360 unknowns (60 photos) was comfortably within what had been estimated to be the maximum that could be accommodated on the available computer. The adjustment was completed in 14 hours of clock time (just under 5 hours of central processor time).

This experience was the author's motivation for investigating a number of reordering algorithms, that are being used in other fields to determine their applicability to photogrammetric networks. This work, not driven by any immediate need, includes only a fraction of the algorithms available; so the reader should be forewarned against expecting an exhaustive analysis.

BACKGROUND

Most conventional photogrammetric networks follow a regular pattern and it is not difficult to find a near optimum ordering by inspection of a coverage diagram; but it is often convenient to be able to rely on the computer to perform this task. With the present trend toward automation in all phases of photogrammetry and cartography, the demand for automating the data ordering process will continually increase. Therefore, it will be valuable to know the characteristics and applicability of some of the automated data ordering algorithms that have become invaluable in many disciplines involving large data adjustment problems.

Duff (1976) defines a sparse matrix system as one in which one can take advantage of either the percentage or distribution of zero elements. The distribution is generally the more important of the two, as evidenced by the advantage that has already been gained with the banded structure of photogrammetric networks. The minimal storage requirements, the small percentage of the inverse elements that have to be computed, and the ease with which logical groups of data can be moved to and from peripheral storage devices create an aesthetically appealing process. Is there room for improvement? Or, should we be satisfied with merely finding a better means of automating the bandwidth minimization process? In the following sections, we will consider the merits of some reordering algorithms used for bandwidth reduction and alternative methods designed to minimize the creation of non-zero elements during matrix factorization; but first some general background information will be needed.

Most sparse matrix literature, and reorder algorithms in particular, rely heavily on graph theory, which like most specialties, has developed a terminology that may not be familiar to the uninitiated. If we let A be an n by n sparse symmetric, positive definite matrix, the diagonal elements, a_{ii}, are called nodes or vertices and the non-zero off-diagonal elements, a_{ij}, are called edges. The maximum value within row i of j-i for non-zero a_{ij} is called the local bandwidth, b_i, of row i; and the maximum of the b_i is called the bandwidth of A. The maximum value of j-i within column j for which a_{ij} is non-zero is called the local column bandwidth; and the sum of all column bandwidths is called the profile (or column profile) of A.

Two nodes a_{ii} and a_{jj} are said to be adjacent if they are connected by an edge a_{ij}. The degree of a node a_{ii} is the number of edges it shares with other nodes or simply the number of non-zero off-diagonal elements in row i.

A path between two nodes is a sequence of edges beginning at one and ending at the other. The distance between two nodes is the length of, or number of edges in, a shortest path from one to the other, and a diameter of the graph is a shortest path connecting two nodes of maximal distance apart. A level structure of a graph is a partition of the nodes into levels such that all nodes adjacent to nodes in level i are in either level i-1, i, or i+1.

In the normal equation matrix associated with a photogrammetric network, the 6 by 6 submatrices lying on the principal diagonal and associated with one photograph can be considered a single node, and the off-diagonal 6 by 6 submatrices as individual edges. This device of treating submatrices as single elements simplifies the analysis considerably. We must keep in mind this difference in terminology, however, when estimating storage or number of operations to be performed. To prevent confusion, we will use B for a bandwidth composed of submatrices and b for one given in terms of matrix elements. To be consistent with previous photogrammetric usage, the bandwidth B will include the diagonal block so that $b = nB - 1$, where n is the dimension of the block used in determining B. Bandwidth is very important because it determines the amount of core storage that must be made available to perform the network adjustment.

Another important factor in sparse matrix methods is the number of elements that are initially zero, but become non-zero as a result of fill during the forward reduction by Gaussian or Cholesky factorization. In photogrammetric networks, fill can be described in terms of a number of matrices, rather than individual elements. Throughout this paper, fill will be denoted by F and will be measured in n by n matrices, where n = 6 unless otherwise specified.

The initial form of the normal equations of a photogrammetric network can be partitioned into two block diagonal matrices, one associated with ground points and the other associated with camera station parameters, and a matrix of connections between them. When ground points are eliminated in the reduction, there is fill that ties together all photographs that image common ground points and is independent of the ordering of the photos. Note that the diagonal blocks were initially non-zero and are the only non-zero matrices that are not a part of F. Once the factorization of the photo parameter partition begins, there will be additional fill that is very much dependent on the ordering of the photos and will determine how many additional inverse terms must be computed. This additional fill, which is a subset of F, will be denoted by F'.

Since the development within the National Ocean Service (NOS) of equipment and techniques for photogrammetric densification of geodetic networks, pioneered by Dr. Schmid and described by Slama (1978), we have become accustomed to thinking of photogrammetric networks as having the same amount of side overlap as forward overlap, usually 67 per cent. Such a configuration has the distinct advantage that, neglecting edge effects, all ground points are imaged on at least nine photographs and all photographs see at least nine ground points. Therefore, a sample case to be used in this evaluation will be a hypothetical network consisting of 6 strips of 8 photos each with the uniform 67 per cent overlap described above and with one ground point located at the nadir of each photo. Photos in the center of the network, therefore, become connected to the twenty-four photos surrounding them. This arrangement removes all asymmetries except for the rectangular form of the network.

BANDWIDTH AND PROFILE REDUCTION ALGORITHMS

The purpose of bandwidth and profile reduction methods is to either minimize the in-core storage requirement by reducing the bandwidth or to minimize computer time by reducing the fill or both. The most widely used method of bandwidth and profile reduction used in photogrammetry is cross-strip numbering done by manual sorting of the data. While this is not an automatic algorithm, it will serve as a model against which other methods can be compared.

CROSS STRIP NUMBERING (CSN).

Cross-strip numbering, as the name implies, is simply numbering the photographs in the order in which they are encountered going across the strips on which they were acquired. This assumes, of course, that the photography was flown in strips parallel to the long dimension of the rectangular ground point network.

For the sample network, the numbering will be from 1 through 6 across the first photos of the strips, 2 through 12 across the next set, etc., as shown in figure 1a. The resulting banded normal equation structure, figure 1b, should be familiar to nearly everyone familiar with photogrammetric bundle adjustment methods. The bandwidth of $B = 2N + 3 = 15$, where N is the number of strips, and additional fill of $F' = 156$ are near optimum.

CUTHILL-MCKEE ALGORITHM

The first step in any bandwidth reduction algorithm is the selection of a starting node. The choice that will lead to the minimum bandwidth is a node of low degree, but not necessarily one of the nodes of minimum degree. In choosing a starting node for the Cuthill-McKee algorithm, two possible upper bounds for the degrees of favorable candidates are suggested: the median degree of all nodes and the mean of the minimum and maximum degrees encountered in the total set. Once chosen, the starting node is assigned number 1, and all nodes adjacent to it are numbered in sequence in order of increasing degree. Ties are broken arbitrarily. Next, all unnumbered nodes adjacent to node 2 are numbered sequentially in order of increasing degree. The first such node is assigned the number following the highest number assigned to the nodes adjacent to node 1. This procedure is then repeated for node 3, node 4, and so forth, until all nodes are numbered.

The numbering of our sample network that results from this scheme is shown in figure 1a. This numbering is not unique, because any of the four corner photographs used as starting node would have produced a equivalent normal equations data structure to that shown in part b of the figure. Furthermore, a number of ties were broken arbitrarily, as specified by the algorithm, and these choices have influenced the structure. In fact, the bandwidth and fill would have increased if certain of these ties had been resolved differently. It is evident that the bandwidth, already greater than that of CSN, will continue to grow as more strips are added, but not if the length of the strips are increased. The bandwidth will be given by $B = 4N - 3$.

REVERSE CUTHILL-MCKEE ALGORITHM (RCM)

George (1971) discovered that reversing the numbering that results from the Cuthill-McKee algorithm will always reduce the profile of the normal equation matrix. Since reversing the order will not change the bandwidth, this Reverse Cuthill-McKee algorithm will nearly always reduce the fill.

Figure 2a shows the reverse of the photo numbering scheme of figure 1a, and the associated data structure (fig. 2b) shows the expected reduction in F' to be impressive--184 as opposed to 254 for conventional Cuthill-McKee. This procedure should certainly be considered to be the more advantageous, but does not compare very favorably to CSN in the sample case.

ALGORITHM OF GIBBS, POOLE, AND STOCKMEYER

Gibbs et al (1974) suggest an algorithm which they claim typically produces bandwidth and profile which are comparable to those of RCM, but accomplishes the reordering in significantly less computation time. A complete description of this algorithm is beyond the scope of this paper. Briefly, the method consists of: 1) finding a pair of nodes that are nearly maximal distance apart by generating level structures; 2) combining the level structures rooted in these two nodes into a new level structure whose width is usually less than either of the original ones; and 3) numbering the nodes within each level of the new structure using a procedure similar to the Cuthill-Mckee algorithm.

When applied to simpler networks in which each point is connected to just the eight neighbors surrounding it, this algorithm provides the ideal cross-strip ordering and gives rise to the hope that this will be an efficient automated procedure for CSN, but for photogrammetric networks the results are disappointing. When applied to the sample case, a bandwidth of B = 20 (one less than RCM) and F = 185 (the same as RCM) result. As promised by its authors the algorithm is a slight improvement over RCM and is obtained in a fraction of the computing time; however, it does not compare favorably with CSN.

BANKER'S ALGORITHM

The banker's algorithm, proposed by Snay (1976), was devised as a means of obtaining near minimal column profile. Once a starting node has been selected, all nodes adjacent to it (its neighbors) are added to a list of hopefuls. All nodes adjacent to either the starting node or hopeful nodes are then added to a list of candidates, unless they are already selected or are already candidates. The next node to be selected from the list of candidates is the one with the minimum number of neighbors that have neither been selected nor added to the hopeful list. In case of ties, nodes on the hopeful list are chosen. Otherwise, ties are broken arbitrarily. When a new node is numbered, all of its neighbors are added to the hopeful list and any of their neighbors not already included are added to the candidate list. This procedure is repeated until all nodes are numbered.

The ordering shown in figure 4a is not unique because of the arbitrary tie breaking procedure employed. The data structure (fig. 4b) obtains F' = 138, a significant improvement over the

156 of CSN, but at the expense of a bandwidth of 27. Both bandwidth and fill would increase if additional strips were added, here B = 4N + 3, but increasing the lenghth of the strips would not affect bandwidth and the increase in F' would be comparable to that with CSN. This appears to be the best algorithm, of those tested, for reducing the profile of a photogrammetric network, but is a poor choice for bandwidth reduction.

A HEURISTIC APPROACH

The author, having searched in vain for an algorithm that would improve upon CSN, or even duplicate its results in the ideal case, did not want to leave the reader with only negative results. If we assume that most networks will by design approximate the ideal case, then any algorithm that will reproduce CSN under ideal circumstances should be worth pursuing. The following algorithm, which has proved to be quite satisfactory in dealing with real networks, was developed in an attempt to construct a set of rules that will cause the computer to produce the CSN ordering:

1) Choose a starting node and label it number 1. Any of the four nodes of minimum degree is obviously a valid choice for the theoretical network under consideration.

2) Form a list of candidates that consists of all neighbors (adjacent nodes) of the starting node.

3) Choose from the the candidate list the node with the fewest neighbors that have not yet become candidates.

a. Ties are broken by choosing the candidate with the maximum number of neighbors that are already on the candidate list.

b. If a tie still exists, choose the candidate whose sponsor (see step 4) has the lowest number.

c. Ties that still exist are to be broken arbitrarily or by a rule, or set of rules, specified by the user (see discussion following step 8).

4) Assign to this node the next number in sequence and add to the candidate list all of its neighbors that are not already candidates. Tag each of these new candidates with the number assigned to the node that caused them to become candidates (their sponsor).

5) Repeat steps 3 and 4 until all neighbors of the starting node have been numbered.

6) Choose from the candidate list the node whose sponsor has the lowest number. In the event of a tie, select the candidate of minimum degree.

7) Follow the procedure given in step 4.

8) Repeat steps 6 and 7 until all nodes are numbered.

For this algorithm to meet its stated objective, to teach the computer to duplicate the manual sorting process, the four equally valid choices for a starting node are known in advance. Having made this choice, however, there does not seem to be a simple means of choosing node number 2 without searching down two paths to find which is longer. There are two candidates for the number 2 position that are indistinguishable through all the tie breaking procedures in step 3. One leads to CSN while the other will produce along-strip numbering.

There are a number of additional tie-breaking procedures that could be applied and the most advantageous and/or efficient choice will depend on the set of circumstances. Some suggestions are:

1) Choose one of the candidates and proceed through step 5. At this point the highest number assigned is the local bandwidth of node 1. Record this number and repeat these steps using the alternative candidate. The candidate that gives the smaller local bandwidth for node 1 is the correct choice.

2) Specify in advance a maximum acceptable bandwidth for the specific network or to apply to all networks, if determined by physical limitation such as storage. Choose one of the candidates and proceed. If a number larger than the specified maximum is reached before step 5 is complete, stop and choose the alternate candidate. If the specified maximum must not be exceeded for some reason, the algorithm should be instructed to stop and take an alternate route if any local bandwidth becomes too large. The growth of the local bandwidth can be monitored at all times by checking the difference between the number assigned to a node and the number of its sponsor.

3) Let the user choose the first two nodes, because he will certainly know which are the along- and across-strip directions, if the network length is much greater than its width.

Without some means of directing its choice of node number 2, this algorithm is equally likely to choose along-strip numbering as CSN, when applied to networks of the type described in this paper. However, when applied to photo networks in which side overlap does not exceed 50 percent, it will always number along the strips.

CONCLUSION

Of the algorithms investigated, the bandwidth reduction methods are the most convenient to employ because most software for block bundle adjustments has been designed for banded matrices. For networks of moderate size, any of the bandwidth methods will suffice, if some inefficiency in computer utilization can be tolerated. All of the tested algorithms need additional tie-breaking rules, however, to perform as well as the idealized applications given in this paper.

If storage is not as important a consideration as speed, the banker's algorithm appears to be a good choice. If most networks are expected to be of the densification type, having side overlap of 60 percent or more, then the heuristic algorithm seems to be a better all-round choice.

ACKNOWLEDGEMENTS

The author is indebted to Dr. Hellmut Schmid whose trust and guidance precipitated many interesting and challenging projects, including the one that led to this investigation. The author also wishes to acknowledge the help, suggestions, and encouragement given him in this investigation at various times over a period of several years by William Dillinger, Robert Hanson, and Allen J. Pope.

REFERENCES

Cuthill, E. and J. Mckee, 1969. Reducing the bandwidth of sparse symmetric matrices, Proceedings of Association for Computing Machinery, 24th National Conference, 157-172.

Doyle, Frederick J., Atef A. Elassal, and James R. Lucas, 1977. Selenocentric geodetic reference system, NOAA Technical Report NOS 70 NGS 5.

Duff, I. S., 1976. A survey of sparse matrix research, Computer Science and Systems Division, AERE Harwell Report.

Elassal, A.A., R. K. Brewer, G. Gracie, and M. Crombie, 1970. MUSAT IV, Final Technical Report, Raytheon Company, Alexandria, Virginia.

George, Alan, 1971. Computer implementation of the finite element method, Report STAN-CS-71-208, Stanford University Computer Science Department, Stanford, California.

Gibbs, Norman E., William G. Poole, Jr., and Paul K. Stockmeyer, 1976. An algorithm for reducing the bandwidth and profile of a sparse matrix, SIAM Journal of Numerical Analysis, 13, pp. 236-250.

Gyer, M., 1967. The inversion of the normal equations of analytical photogrammetry by the method of recursive partitioning, RADC TR-67-69, Rome Air Development Center, Rome, New York.

Slama, C. C., 1978. High precision analytical photogrammetry using a special reseau geodetic lens cone. Proceedings of ISP Commission III International Symposium, International Society of Photogrammetry, 31 July-5 August, Moscow, U.S.S.R.

Snay, R. A., 1976. Reducing the profile of sparse symmetric matrices, Bulletin Geodesique, 50, pp. 341-352.

FIGURE CAPTIONS

1	7	13	19	25	31	37	43
2	8	14	20	26	32	38	44
3	9	15	21	27	33	39	45
4	10	16	22	28	34	40	46
5	11	17	23	29	35	41	47
6	12	18	24	30	36	42	48

(a) Numbering of Photographs

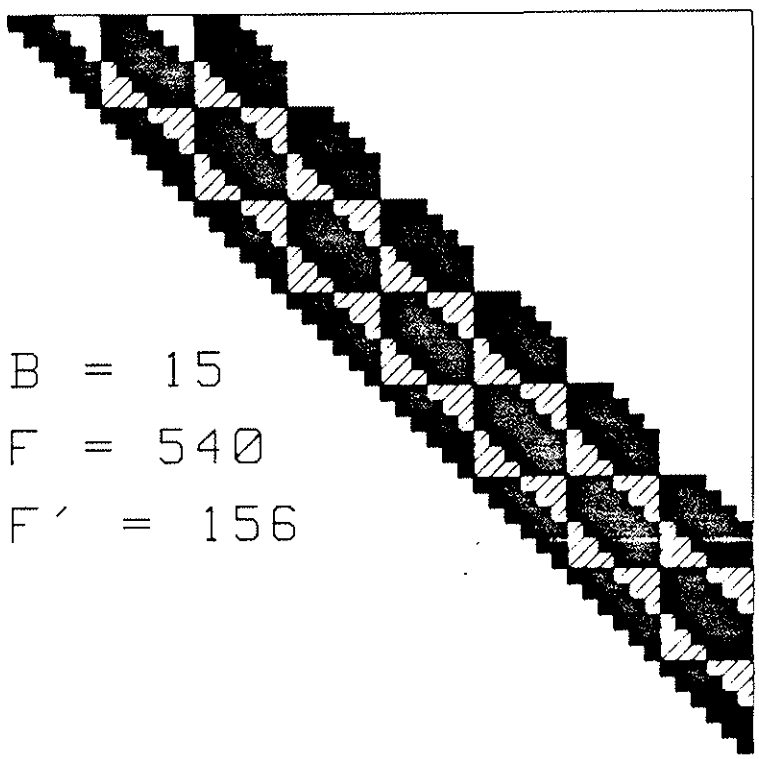

(b) Normal Equation Structure

48	46	44	36	30	18	13	6
47	43	41	35	29	17	12	5
45	42	40	34	28	16	11	4
39	38	37	27	25	15	10	3
33	32	31	26	24	14	9	2
23	22	21	20	19	8	7	1

(a) Numbering of Photographs

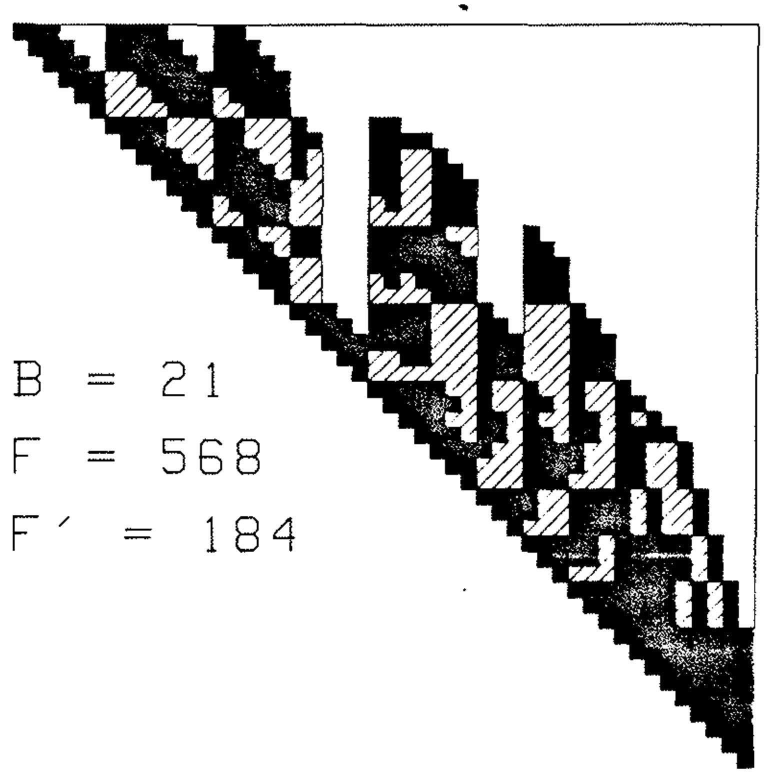

(b) Normal Equation Structure

1	7	13	19	25	33	32	31
2	8	14	20	26	36	35	34
3	9	15	21	27	39	38	37
4	10	16	22	28	42	41	40
5	11	17	23	29	45	44	43
6	12	18	24	30	48	47	46

(a) Numbering of Photographs

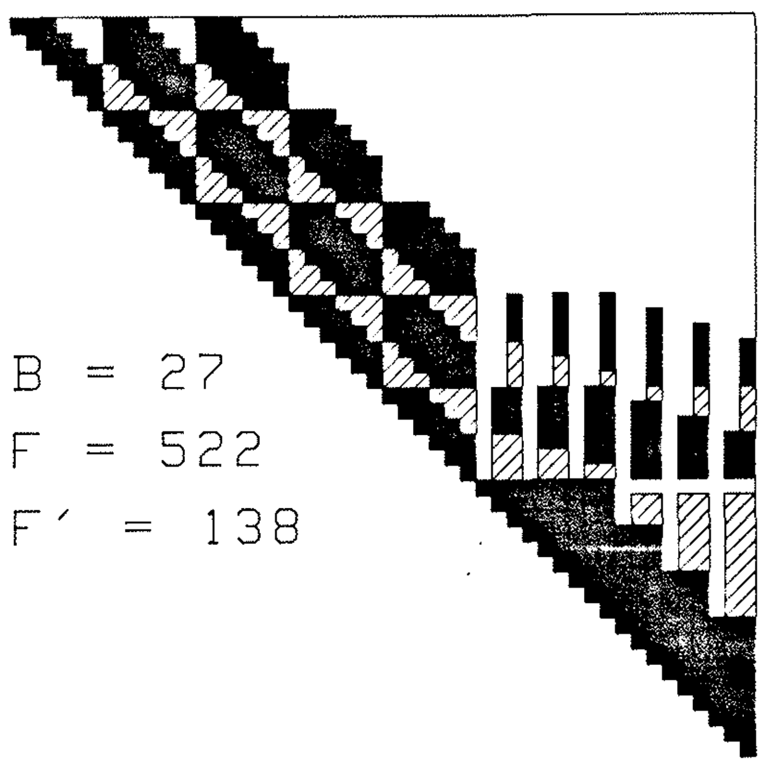

(b) Normal Equation Structure