



Algorithms

Submission of an algorithm for consideration for publication in Communications of the ACM implies unrestricted use of the algorithm within a computer is permissible.

Copyright © 1973, Association for Computing Machinery, Inc. General permission to republish, but not for profit, all or part of this material is granted provided that ACM's copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery.

Algorithm 455

Analysis of Skew Representations of the Symmetric Group [Z]

D.B. Hunter* and Julia M. Williams† [Recd. 5 Feb. 1971]

* Department of Mathematics, University of Bradford, Yorkshire, England

† 12 Peel Close, Heslington, York, England

Key Words and Phrases: symmetric group, skew representation, partition, Young diagram, lattice permutation, binary model, outer product

CR Categories: 5.30

Language: Algol

Description

This algorithm analyzes the skew representation $[\lambda]-[\mu]$ of the symmetric group σ_n corresponding to a pair of partitions

$(\lambda) = (\lambda_1, \lambda_2, \dots, \lambda_r)$ and $(\mu) = (\mu_1, \mu_2, \dots, \mu_s)$ where

$$\left. \begin{aligned} r &\geq s \\ \lambda_1 &\geq \lambda_2 \geq \dots \geq \lambda_r \\ \mu_1 &\geq \mu_2 \geq \dots \geq \mu_s \\ \lambda_i &\geq \mu_i \quad (1 \leq i \leq s) \\ n &= \sum_{i=1}^r \lambda_i - \sum_{i=1}^s \mu_i \end{aligned} \right\} \quad (1)$$

(see Robinson [4, sec. 2.5]). The analysis takes the form

$$[\lambda]-[\mu] = \sum_{(\nu)} c_{(\nu)} [\nu], \quad (2)$$

where the summation is over all partitions (ν) of n , the coefficients $c_{(\nu)}$ being nonnegative integers.

The method used may be described as follows: construct all possible diagrams which can be built up in accordance with the following two rules.

(a) Replace μ_s of the nodes in the Young diagram corresponding to (λ) by identical symbols α_s in such a way that: (i) the unchanged

nodes form a regular Young diagram; and (ii) no two identical symbols α_s lie in the same column. Then replace μ_{s-1} further nodes by identical symbols α_{s-1} in accordance with the same rules, and so on, finally replacing μ_1 nodes by identical symbols α_1 .

(b) In the final diagram the altered nodes should form a lattice permutation of $\alpha_1^{\mu_1} \alpha_2^{\mu_2} \dots \alpha_s^{\mu_s}$ (Robinson [4, sec. 2.4]) when read from right-to-left through successive rows.

Then the pattern of unchanged nodes in each diagram so constructed defines a term $[\nu]$ in the analysis.

This method appears not to have been explicitly stated in the above form before, but is an immediate consequence of Littlewood's method for analyzing the outer product $[\lambda] \cdot [\mu]$ (see Littlewood [3, sec. 6.3, th. V], Robinson [4, sec. 3.3]), noting that $c_{(\nu)}$ is also the coefficient of $[\lambda]$ in the analysis of $[\mu] \cdot [\nu]$ (Littlewood [3, sec. 6.4, th. VIII]).

In the procedure, binary models of those partitions (ν) in (2) for which $c_{(\nu)} \neq 0$ are stored, in lexicographic order, in $nu[1], nu[2], \dots, nu[p]$, the corresponding values $c_{(\nu)}$ being stored in $c[1], c[2], \dots, c[p]$. The binary model used is due to Com et [1], a partition $(\nu) = (\nu_1, \nu_2, \dots, \nu_t)$ being represented by the number $2^{n-\nu_1} + 2^{n-\nu_1-\nu_2} + \dots + 2^{\nu_t} + 1$. (3)

The techniques used are similar to those employed in [2]. In particular, two two-dimensional arrays *lam* and *sigma* are required. Corresponding to any particular diagram, *lam* $[i, j]$ specifies the number of nodes in row j which are still unchanged when all the symbols $\alpha_s, \alpha_{s-1}, \dots, \alpha_i$ have been inserted ($j = i, i + 1, \dots, r$), and *sigma* $[i, j]$ specifies the total number of symbols α_i inserted in rows $i, i + 1, \dots, j$. Thus the quantities *lam* $[i, j]$ are generated by the equation

$$lam[i, j] = lam[i+1, j] - sigma[i, j] + sigma[i, j-1]. \quad (4)$$

The rules for constructing the diagrams impose the restrictions

$$sigma[i-1, j-1] \geq sigma[i-1, j] - lam[i, j] + lam[i, j+1] \quad (5)$$

and

$$sigma[i-1, j-1] \geq sigma[i, j]. \quad (6)$$

Each time array *lam* is completed, a term

$$(\nu) = (lam[1, 1], lam[1, 2], \dots, lam[1, r]) \quad (7)$$

is added to the analysis.

Note 1. In view of the identity

$$[\lambda] \cdot [\mu] = [\lambda_1 + \mu_1, \lambda_1 + \mu_2, \dots, \lambda_1 + \mu_s, \lambda_1, \lambda_2, \dots, \lambda_r] - [\lambda_1^s],$$

procedure *skew* may also be used to analyse the outer product $[\lambda] \cdot [\mu]$. It is, however, less convenient for this purpose than procedure *outer product* of Hunter [2].

Note 2. Value of p . It is difficult to predict the value of p in any example. Clearly, $p \leq p(n)$, where $p(n)$ denotes the number of partitions of n . On the other hand, for any value of n , there are partitions (λ) and (μ) for which $p = p(n)$, namely, $(\lambda) = (n, n-1, \dots, 1)$, $(\mu) = (n-1, \dots, 1)$.

References

- Com et, S. Notations for partitions. *MTAC* 9 (1955), 143-146.
- Hunter, D.B. Outer product of symmetric group representations. *BIT* 10 (1970), 106-114.
- Littlewood, D.E. *Theory of Group Characters*, 2nd ed. Oxford U. Press, England, 1950.
- Robinson, Gilbert B. *Representation Theory of the Symmetric Group*. U. of Toronto Press, Toronto, Ont., Canada, 1961.

Algorithm

procedure *skew* ($r, s, lambda, mu, p, c, nu$);

value r, s ; **integer** r, s, p ; **integer array** *lambda, mu, c, nu*;

begin

comment Input parameters.

r : the number of parts in partition (λ) .

s : the number of parts in partition (μ) .

If, in the middle of tour building, the tour, for instance, consists of the nodes p_1, p_2, \dots, p_k , the next node among the nodes $p_{k+1}, p_{k+2}, \dots, p_n$, and the arc (to be split by the chosen inserted node) among the arcs $p_1p_2, p_2p_3, \dots, p_kp_1$, are chosen so that the tour increment will be minimum; i.e. i ($1 \leq i \leq k$) and j ($k < j \leq n$) are chosen in such a manner that $d_{(p_i, p_j)} + d_{(p_j, p_{i+1})} - d_{(p_i, p_{i+1})} = \min$. Tour building starts with the arc p_1p_1 and terminates when all n nodes have been included. The tour-building approach of this kind for the traveling-salesman problem was originated by Karg and Thompson [1] and further developed by Raymond [2]. This algorithm, however, handles an open connection—start-to-end or start-to-any node. The maintenance of this property is ensured in the algorithm by assigning to the end-to-start or each-to-start distance sufficiently large negative values ($-n \times \max_{ij} [d_{ij}]$) which, in some way, firmly attach the end or any of n nodes to the start node permitting a circuit to form. In fact, the algorithm works on a net as if it were a closed circuit and keeps the node configuration by modifying the distance matrix. In output, the distance matrix is returned to its original form.

A tour thus built is hardly optimal and for larger nets it is probably far from optimum. The second phase improves the tour (for $n \geq 3$) by the so-called 3-opt method proposed by Lin [3]. Improvements consist in exchanging three arcs, or links of the given connection by three other links. If there are no more 3 links to exchange for tour improvement, the tour is said to be 3-optimal. In general, λ -optimality can be considered. The implication of the 3-link exchange is essentially in reinsertions. Consecutive node chains of length k ($1 \leq k < n$) are successively tried to be reinserted (both as are and inverted) into remaining links for tour improvements, which actually represent 3-link exchanges (and also 2-link at the same time). A 3-opt tour shows a certain probability to be an optimal one in relation to n . Different 3-opt tours can be achieved if different initial nodes are chosen, which allows us to increase the probability of obtaining an optimal solution.

The algorithm can run r trials (as specified in input) with different initial nodes (p_1 , set automatically), thus obtaining different solutions while the best is saved and replaced in the vector p in output. For runs $r > n$ ($r \leq 2n$) there is little chance for further improvement, because initial nodes repeat and the tour development can be affected only by previous contents of the vector p on which the tour is built. Probability that the 3-opt tour is optimal is somewhat higher in this algorithm, than in the one Lin suggests. In contrast to finding a 3-opt solution from a given random sequence of nodes, the fast building of an appropriate tour in the first phase considerably reduces the number of reinsertions in the second phase. The algorithm generalization to the noncyclic and nonsymmetric problems, in comparison to the traveling-salesman problem, increases computational time.

A considerable number of test examples have been run by the algorithm including the three problem types mentioned and the non-Euclidean and nonsymmetric problems. To outline the capability and how the "cost-approximation" factor r should be set for various n 's, a survey of tested problems is presented, most of which problems have been solved and published before. The algorithm in Fortran was run on the RCA's SPECTRA 70/45 (fixed-point add time equals 8.88 μ sec), and is recommended for a high probability (over 95 percent) of obtaining an optimum if $r = 2$ to 5 for $n \leq 10$ and $r = 5$ to 15 for $n \leq 30$. For higher n 's, unless cost is out of consideration and r can be set up to $2n$, the checking of successive results is advisable to see how improvements are developing (p and l_1 should be checked after the tour-length calculation). These checks can also serve for getting suboptimal solutions.

In the program, the distance matrix d is in fixed-point mode, which makes computation faster and does not seem to be a serious restriction. Decimal order range of distances is expected to be small enough to be represented in fixed point, and calculations (additions and subtractions) will, most likely, not face overflow problem.

The arrays ID and Q should have the maximum subscript set at least to n .

Survey of tested problems

Ref.	n	sn	en	Conjectured optimum	r_{opt}	t_1 [sec]	
Karg and Thompson [1]	5	1	2	118	1	<1	
		1	0	108	1		
		1	1	148	1		
Raymond [2]	7	1	5	165	1	<1	
		1	0	140	1		
		1	1	179	1		
Barachet [4]	10	1	2	350	1	1.4	
		1	0	298	1		
		1	1	378	2		
		1	2	308	1		
		10*	1	0	257		2
Author	12	1	2	102	1	3.0	
		1	0	95	1		
		1	1	114	1		
Author	13	1	6	117	1	3.0	
		1	0	102	1		
		1	1	130	1		
Held and Karp [5]	25	1	25	**	1517	10	21.8
		1	0	** $en=25$	1517	2	22.3
		1	1	**	1711	1	29.7
Karg and Thompson [1]	33	1	33	**	10655	2	53.6
		1	0	** $en=14$	10585	10	53.4
		1	1	**	10861	6	53.7

* Nonsymmetric problem (two distances changed: $(6, 5) = 1$, and $(8, 3) = 1$).
** Results obtained from 10 runs.

The algorithm is believed to be applicable also to problems in which all connections do not necessarily exist. In terms of graph theory a graph representing the net to be routed need not be complete; i.e. every pair of vertices may be connected only in one of the two possible directions. The graph, however, must be strongly connected; i.e. there must be a path joining any pair of arbitrary distinct vertices. Nonexisting arcs might be expressed by assigning to the appropriate distances d_{ki} sufficiently large positive values, for instance $n \times \max_{ij} [d_{ij}]$.

Symbol summary

- n number of nodes to be connected ($2 \leq n \leq m$).
- p vector containing n node numbers (in output, it contains node number sequence of conjectured shortest path).
- sn start node number ($1 \leq sn \leq m$; no check is provided whether sn is contained in p).
- en end node number ($1 \leq en \leq m$; if $en = 0$, start-to-any connection is assumed; $en = sn$ is allowed, which is traveling-salesman problem; no check is provided whether sn is contained in p).
- m order of distance matrix d ($m \geq n \geq 2$).
- d $m \times m$ matrix of distances of all node pairs (zero diagonal, not necessarily symmetric).
- l length of conjectured shortest path (output).
- r number of runs (trials; $r \leq 2n$).
- r_{opt} serial run number during which optimum has been achieved.
- t_1 average computational time of one run in seconds.

References

1. Karg, R.L., and Thompson, G.L. A heuristic approach to solving traveling salesman problem. *Mgmt. Sci.* 10, 2 (1964), 225-248.
2. Raymond, T.C. Heuristic algorithm for the traveling-salesman problem. *IBM J. Res. Develop.* 13, 4 (1969), 400-407.

3. Lin, S. Computer solutions of the traveling salesman problem. *Bell Syst. Tech. J.* 44 (Dec. 1965), 2245-2269.
4. Barachet, L.L. Graphic solution of the traveling salesman problem. *Oper. Res.* 5 (1957), 841-845.
5. Held, M., and Karp, R.M. A dynamic programming approach to sequencing problems. *J. Soc. Indust. Appl. Math.* 10, 1 (1962), 196-210.
6. Saksena, J.P., and Kumar S. The routing problem with 'k' specified nodes. *Oper. Res.* 14 (1969), 909-913.
7. Bellmore, M., and Nemhauser, G.L. The traveling salesman problem: A survey. *Oper. Res.* 16 (1968), 538-558.
8. Berge, C. *The Theory of Graphs and Its Applications*. Wiley, New York, 1962.
9. Berge, C., and Ghouila-Houri, A. *Programming, Games and Transportation Networks*. Wiley, New York, 1965.

Algorithm

```

SUBROUTINE ROUTNG(N, P, SN, EN, M, D, L, R)
INTEGER P(N), D(M,M), ID(60), O(60), SN, EN, R
C N - NUMBER OF NODES TO BE CONNECTED
C P - NODE NUMBER VECTOR (IN OUTPUT, OPTIMAL CONNECTION)
C SN- START NODE NUMBER
C EN- END NODE NUMBER
C M - DISTANCE MATRIX ORDER
C D - DISTANCE MATRIX
C L - SHORTEST CONNECTION LENGTH (OUTPUT)
C R - NUMBER OF RUNS
C GET LARGE NUMBER (= N X MAX D(I,J))
LARGE = 0
DO 20 I=1,M
DO 10 J=1,M
IF (D(I,J).GT.LARGE) LARGE = D(I,J)
10 CONTINUE
20 CONTINUE
LARGE = LARGE*N
C DEFINE NON-EXISTING ARCS BY ASSIGNING
C THEIR DISTANCES LARGE NEGATIVE VALUES
IF (EN.EQ.0) GO TO 40
DO 30 I=1,M
ID(I) = D(I,SN)
D(I,SN) = -LARGE
D(SN,SN) = 0
30 CONTINUE
40 IF (SN.EQ.EN .OR. EN.EQ.0) GO TO 50
ID(1) = D(EN,SN)
D(EN,SN) = -LARGE
C RUN R TRIALS
50 L = LARGE
DO 280 IRS=1,R
C BUILD TOUR BY SUCCESSIVE INSERTING
C NOT-YET-INVOLVED NODES
C INITIATE TOUR IS CONSIDERED AS
C ARC P(1) TO P(1)
DO 90 JS=2,N
MININC = LARGE
C TRACE ALL NOT-YET-INVOLVED NODES
C TO CHOOSE THE ONE WITH MINIMUM INCREMENT
DO 70 J=JS,N
JP = P(J)
JE = JS - 1
C FOR EACH NOT-YET-INVOLVED NODE TRACE ALREADY
C BUILT-UP TOUR TO CHOOSE THE MINIMUM INCREMENT ARC
DO 60 I=1,JE
IP = P(I)
IP1 = P(I+1)
IF (I.EQ.JE) IP1 = P(1)
INC = D(IP,JP) + D(JP,IP1) - D(IP,IP1)
IF (INC.GE.MININC) GO TO 60
J1 = J
I1 = I
MININC = INC
60 CONTINUE
70 CONTINUE
C STRETCH TOUR BY INSERTING THE CHOSEN NODE P(J1)
C BETWEEN THE NODES P(I1) AND P(I1+1)
80 J1 = J1 - 1
IF (J1.EQ.I1) GO TO 90
IP = P(J1)
P(J1) = P(J1+1)
P(J1+1) = IP
GO TO 80
90 CONTINUE
C CORRECT TOUR BY 3-ØPT METHOD
C VARY CONSECUTIVE CHAIN LENGTH K
N1 = N - 1
IF (N.LT.3) GO TO 210
DO 200 K=1,N1
ICOUNT = 0
C SHIFT CONSECUTIVE CHAIN
C THROUGHOUT SEQUENCE OF N NODES
100 ICØR = 0
DO 190 J=1,N
C CALCULATE CHAIN LENGTH IN FORWARD
C AND BACKWARD DIRECTION
L1 = 0
LR = 0
IF (K.EQ.1) GO TO 120
I = J
K1 = 1
IF (I.GT.N) I = I - N
IP = P(I)
IP1 = I + 1
IF (IP1.GT.N) IP1 = 1
IP1 = P(IP1)

```

```

L1 = L1 + D(IP,IP1)
LR = LR + D(IP1,IP)
I = I + 1
K1 = K1 + 1
IF (K1.LT.K) GO TO 110
C FOR EACH POSITIONED CHAIN (AS IS AND INVERTED)
C CHECK ALL ARCS IF INSERTION IMPROVES TOUR
120 MININC = LARGE
J1 = J + K - 1
IF (J1.GT.N) J1 = J1 - N
DO 150 I=1,N
IF (J.LE.J1 .AND. (I.GE.J .AND. I.LE.J1)) GO TO 150
* IF (J.GT.J1 .AND. (I.LE.J1 .OR. I.GE.J)) GO TO 150
IP = P(I)
JP = P(J)
JP1 = P(J1)
IP1 = I + 1
IF (IP1.GT.N) IP1 = 1
JE = IP1
IF (IP1.EQ.J) IP1 = J1 + 1
IF (IP1.GT.N) IP1 = 1
IP1 = P(IP1)
LN = L1
IR = 0
130 INC = D(IP,JP) + LN + D(JP1,IP1) - D(IP,IP1)
IF (INC.GT.MININC .OR. (INC.EQ.MININC .AND.
(JE.NE.J .OR. JE.EQ.J .AND. IR.EQ.1)) GO TO 140
* I1 = I
IR1 = IR
MININC = INC
140 IF (IR.EQ.1) GO TO 150
IR = 1
LN = LR
JS = JP
JP = JP1
JP1 = JS
GO TO 130
150 CONTINUE
I = I1 + 1
IF (I.GT.N) I = 1
IF (I.EQ.J .AND. IR1.EQ.0) GO TO 190
C REINSERT CHAIN OF LENGTH K STARTING IN J
C BETWEEN NODES P(I1) AND P(I1+1)
ICØR = 1
JS = J
JE = 0
IF (IR1.EQ.0) GO TO 160
JS = J1
JE = -1
160 K1 = 0
170 K1 = K1 + 1
IF (K1.GT.K) GO TO 190
I = JS
JS = JS + JE
IF (JS.LT.1) JS = N
IP = I + 1
IF (IP.GT.N) IP = 1
JP = P(I)
P(I) = P(IP)
P(IP) = JP
I = I + 1
IF (I.GT.N) I = 1
IF (IP=I) 180, 170, 180
180 CONTINUE
IF (ICØR.EQ.0) GO TO 200
ICOUNT = ICOUNT + 1
IF (ICOUNT.LT.N) GO TO 100
200 CONTINUE
C ORIENT TOUR WITH SN IN P(1)
210 DO 230 I=1,N
IF (P(I).EQ.SN) GO TO 240
JS = P(1)
DO 220 J=1,N1
P(J) = P(J+1)
220 CONTINUE
P(N) = JS
230 CONTINUE
C CALCULATE TOUR LENGTH
240 L1 = 0
DO 250 I=1,N1
IP = P(I)
IP1 = P(I+1)
L1 = L1 + D(IP,IP1)
250 CONTINUE
IP = P(I)
IF (SN.EQ.EN) L1 = L1 + D(IP1,IP)
C SAVE SOLUTION, IF BETTER, AND SET NEW INITIATE NODE
IF (L1.GE.L) GO TO 270
L = L1
DO 260 I=1,N
O(I) = P(I)
260 CONTINUE
270 J = IRS + 1
IF (J.GT.N) J = J - N
JS = P(1)
P(1) = P(J)
P(J) = JS
280 CONTINUE
C RESTORE P AND DUMMY DISTANCES
DO 290 I=1,N
P(I) = O(I)
290 CONTINUE
IF (EN.EQ.0) GO TO 310
DO 300 I=1,N
O(I,SN) = ID(I)
300 CONTINUE
310 IF (SN.EQ.EN .OR. EN.EQ.0) GO TO 320
D(EN,SN) = ID(1)
320 RETURN
END

```

Algorithm 457

Finding All Cliques of an Undirected Graph [H]

Coen Bron* and Joep Kerbosch† [Recd. 27 April 1971 and 23 August 1971]

* Department of Mathematics † Department of Industrial Engineering, Technological University Eindhoven, P.O. Box 513, Eindhoven, The Netherlands

Present address of C. Bron: Department of Electrical Engineering, Twente University of Technology, P.O. Box 217, Enschede, The Netherlands.

Key Words and Phrases: cliques, maximal complete subgraphs, clusters, backtracking algorithm, branch and bound technique, recursion

CR Categories: 3.71, 5.32

Language: Algol

Description

Introduction. A maximal complete subgraph (clique) is a complete subgraph that is not contained in any other complete subgraph.

A recent paper [1] describes a number of techniques to find maximal complete subgraphs of a given undirected graph. In this paper, we present two backtracking algorithms, using a branch-and-bound technique [4] to cut off branches that cannot lead to a clique.

The first version is a straightforward implementation of the basic algorithm. It is mainly presented to illustrate the method used. This version generates cliques in alphabetic (lexicographic) order.

The second version is derived from the first and generates cliques in a rather unpredictable order in an attempt to minimize the number of branches to be traversed. This version tends to produce the larger cliques first and to generate sequentially cliques having a large common intersection. The detailed algorithm for version 2 is presented here.

Description of the algorithm—Version 1. Three sets play an important role in the algorithm. (1) The set *compsub* is the set to be extended by a new point or shrunk by one point on traveling along a branch of the backtracking tree. The points that are eligible to extend *compsub*, i.e. that are connected to all points in *compsub*, are collected recursively in the remaining two sets. (2) The set *candidates* is the set of all points that will in due time serve as an extension to the present configuration of *compsub*. (3) The set *not* is the set of all points that have at an earlier stage already served as an extension of the present configuration of *compsub* and are now explicitly excluded. The reason for maintaining this set *not* will soon be made clear.

The core of the algorithm consists of a recursively defined extension operator that will be applied to the three sets just described. It has the duty to generate all extensions of the given configuration of *compsub* that it can make with the given set of candidates and that do not contain any of the points in *not*. To put it differently: all extensions of *compsub* containing any point in *not* have already been generated. The basic mechanism now consists of the following five steps:

- Step 1. Selection of a candidate.
- Step 2. Adding the selected candidate to *compsub*.
- Step 3. Creating new sets *candidates* and *not* from the old sets by

removing all points not connected to the selected candidate (to remain consistent with the definition), keeping the old sets in tact.

Step 4. Calling the extension operator to operate on the sets just formed.

Step 5. Upon return, removal of the selected candidate from *compsub* and its addition to the old set *not*.

We will now motivate the extra labor involved in maintaining the sets *not*. A necessary condition for having created a clique is that the set *candidates* be empty; otherwise *compsub* could still be extended. This condition, however, is not sufficient, because if now *not* is nonempty, we know from the definition of *not* that the present configuration of *compsub* has already been contained in another configuration and is therefore not maximal. We may now state that *compsub* is a clique as soon as both *not* and *candidates* are empty.

If at some stage *not* contains a point connected to all points in *candidates*, we can predict that further extensions (further selection of candidates) will never lead to the removal (in Step 3) of that particular point from subsequent configurations of *not* and, therefore, not to a clique. This is the branch and bound method which enables us to detect in an early stage branches of the backtracking tree that do not lead to successful endpoints.

A few more remarks about the implementation of the algorithm seem in place. The set *compsub* behaves like a stack and can be maintained and updated in the form of a global array. The sets *candidates* and *not* are handed to the extensions operator as a parameter. The operator then declares a local array, in which the new sets are built up, that will be handed to the inner call. Both sets are stored in a single one-dimensional array with the following layout:

<i>not</i>	<i>candidates</i>
index values: 1..... <i>ne</i> <i>ce</i>

The following properties obviously hold:

- 1. $ne \leq ce$
- 2. $ne = ce$: empty (*candidates*)
- 3. $ne = 0$: empty (*not*)
- 4. $ce = 0$: empty (*not*) and empty (*candidates*)
= clique found

If the selected candidate is in array position $ne + 1$, then the second part of Step 5 is implemented as $ne := ne + 1$.

In version 1 we use element $ne + 1$ as selected candidate. This strategy never gives rise to internal shuffling, and thus all cliques are generated in a lexicographic ordering according to the initial ordering of the candidates (all points) in the outer call.


For an implementation of version 1 we refer to [3].

Description of the algorithm—Version 2. This version does not select the candidate in position $ne + 1$, but a well-chosen candidate from position, say *s*. In order to be able to complete Step 5 as simply as described above, elements *s* and $ne + 1$ will be interchanged as soon as selection has taken place. This interchange does not affect the set *candidates* since there is not implicit ordering. The selection does affect, however, the order in which the cliques are eventually generated.

Now what do we mean by “well chosen”? The object we have in mind is to minimize the number of repetitions of Steps 1–5 inside the extension operator. The repetitions terminate as soon as the bound condition is reached. We recall that this condition is formulated as: there exists a point in *not* connected to all points in *candidates*. We would like the existence of such a point to come about at the earliest possible stage.

Let us assume that with every point in *not* is associated a counter, counting the number of candidates that this point is not connected to (number of disconnections). Moving a selected candidate into *not* (this occurs after extension) decreases by one all counters of the points in *not* to which it is disconnected and introduces a new counter of its own. Note that no counter is ever

Fig. 1. Random graphs show the computing time per clique (in ms) versus dimension of the graph (in brackets: total number of cliques in the test sample).



decreased by more than one at any one instant. Whenever a counter goes to zero the bound condition has been reached.

Now let us fix one particular point in *not*. If we keep selecting candidates disconnected to this fixed point, the counter of the fixed point will be decreased by one at every repetition. No other counter can go down more rapidly. If, to begin with, the fixed point has the lowest counter, no other counter can reach zero sooner, as long as the counters for points newly added to *not* cannot be smaller. We see to this requirement upon entry into the extension operator, where the fixed point is taken either from *not* or from the original *candidates*, whichever point yields the lowest counter value after the first addition to *not*. From that moment on we only keep track of this one counter, decreasing it for every next selection, since we will only select disconnected points.

The Algol 60 implementation of this version is given below.

Discussion of comparative tests. Augustson and Minker [1] have evaluated a number of clique finding techniques and report an algorithm by Bierstone [2] as being the most efficient one.


In order to evaluate the performance of the new algorithms, we implemented the Bierstone algorithm¹ and ran the three algorithms on two rather different testcases under the Algol system for the EL-X8.

For our first testcase we considered random graphs ranging in dimension from 10 to 50 nodes. For each dimension we generated a collection of graphs where the percentage of edges took on the following values: 10, 30, 50, 70, 90, 95. The cpu time per clique for each dimension was averaged over such a collection. The results are graphically represented in Figure 1.

The detailed figures [3] showed the Bierstone algorithm to be of slight advantage in the case of small graphs containing a small number of relatively large cliques. The most striking feature, however, appears to be that the time/clique for version 2 is hardly dependent on the size of the graph.

¹ Bierstone's algorithm as reported in [1] contained an error. In our implementation the error was corrected. The error was independently found by Mulligan and Corneil at the University of Toronto, and reported in [6].

Fig. 2. Moon-Moser graphs show the computing time (in ms) versus k . Dimension of the graph = $3k$. Plotted on logarithmic scale.



The difference between version 1 and "Bierstone" is not so striking and may be due to the particular Algol implementation. It should be borne in mind that the sets of nodes as they appear in the Bierstone algorithm were coded as one-word binary vectors, and that a sudden increase in processing time will take place when the input graph is too large for "one-word representation" of its subgraphs.

The second testcase was suggested by the referee and consisted of regular graphs of dimensions $3 \times k$. These graphs are constructed as the complement of k disjoint 3-cliques. Such graphs contain 3^k cliques and are proved by Moon and Moser [5] to contain the largest number of cliques per node.

In Figure 2 a logarithmic plot of computing time versus k is presented. We see that both version 1 and version 2 perform significantly better than Bierstone's algorithm. The processing time for version 1 is proportional to 4^k , and for version 2 it is proportional to $(3.14)^k$ where 3^k is the theoretical limit.

Another aspect to be taken into account when comparing algorithms is their storage requirements. The new algorithms presented in this paper will need at most $\frac{1}{2}M(M+3)$ storage locations to contain arrays of (small) integers where M is the size of largest connected component in the input graph. In practice this limit will only be approached if the input graph is an almost complete graph. The Bierstone algorithm requires a rather unpredictable amount of store, dependent on the number of cliques that will be generated. This number may be quite large, even for moderate dimensions, as the Moon-Moser graphs show.

Finally it should be pointed out that Bierstone's algorithm does not report isolated points as cliques, whereas the new algorithm does. Either algorithm can, however, be modified to produce results equivalent to the other. Suppression of 1-cliques in the new algorithm is the simplest adaption.

Acknowledgments. The authors are indebted to H.J. Schell for preparation of the test programs and collection of performance statistics. Acknowledgments are also due to the referees for their valuable suggestions.

References

1. Augustson, J.G., and Minker, J. An analysis of some graph theoretical cluster techniques, *J. ACM* 17 (1970), 571–588.
2. Bierstone, E. Unpublished report. U of Toronto.
3. Bron, C., Kerbosch, J.A.G.M., and Schell, H.J. Finding cliques in an undirected graph. Tech. Rep. Technological U. of Eindhoven, The Netherlands.
4. Little, John D.C., et al. An algorithm for the traveling salesman problem. *Oper. Res.* 11 (1963), 972–989.
5. Moon, J.W., and Moser, L. On cliques in graphs. *Israel J. Math.* 3 (1965), 23–28.
6. Mulligan, G.D., and Corneil, D.G. Corrections to Bierstone's algorithm for generating cliques. *J. ACM* 19 (Apr. 1972), 244–247.

Algorithm

```

procedure output maximal complete subgraphs 2(connected, N);
  value N; integer N;
  Boolean array connected;
comment The input graph is expected in the form of a symmetrical
  Boolean matrix connected. N is the number of nodes in the
  graph. The values of the diagonal elements should be true;
begin
  integer array ALL, compsub[1 : N];
  integer c;
  procedure extend version 2(old, ne, ce);
    value ne, ce; integer ne, ce;
    integer array old;
  begin
    integer array new[1 : ce];
    integer nod, fixp;
    integer newne, newce, i, j, count, pos, p, s, sel, minmod;
    comment The latter set of integers is local in scope but need
    not be declared recursively;
    minmod := ce; i := nod := 0;
  DETERMINE EACH COUNTER VALUE AND LOOK FOR
  MINIMUM:
    for i := i + 1 while i ≤ ce ∧ minmod ≠ 0 do
      begin
        p := old[i]; count := 0; j := ne;
  COUNT DISCONNECTIONS:
        for j := j + 1 while j ≤ ce ∧ count < minmod do
          if ¬ connected[p, old[j]] then
            begin
              count := count + 1;
  SAVE POSITION OF POTENTIAL CANDIDATE:
              pos := j
            end;
  TEST NEW MINIMUM:
            if count < minmod then
              begin
                fixp := p; minmod := count;
                if i ≤ ne then s := pos
              else
                begin s := i; PREINCR: nod := 1 end
              end NEW MINIMUM;
            end i;
    comment If fixed point initially chosen from candidates then
    number of disconnections will be preincreased by one;
  BACKTRACKCYCLE:
    for nod := minmod + nod step −1 until 1 do
      begin
  INTERCHANGE:
        p := old[s]; old[s] := old[ne + 1];
        sel := old[ne + 1] := p;
  FILL NEW SET not:
        newne := i := 0;
        for i := i + 1 while i ≤ ne do
          if connected[sel, old[i]] then

```

```

          begin newne := newne + 1; new[newne] := old[i] end;
  FILL NEW SET cand:
          newce := newne; i := ne + 1;
          for i := i + 1 while i ≤ ce do
            if connected[sel, old[i]] then
              begin newce := newce + 1; new[newce] := old[i] end;
  ADD TO compsub:
          c := c + 1; compsub[c] := sel;
          if newce = 0 then
            begin
              integer loc;
              oustring(1, 'clique = ');
              for loc := 1 step 1 until c do
                outinteger(1, compsub[loc])
              end output of clique
            else
              if newne < newce then extend version 2(new, newne, newce);
  REMOVE FROM compsub:
              c := c − 1;
  ADD TO not:
              ne := ne + 1;
              if nod > 1 then
                begin
  SELECT A CANDIDATE DISCONNECTED TO THE FIXED
  POINT:
                  s := ne;
  LOOK: FOR CANDIDATE:
                  s := s + 1;
                  if connected[fixp, old[s]] then go to LOOK
                end selection
              end BACKTRACKCYCLE
            end extend version 2;
            for c := 1 step 1 until N do ALL[c] := c;
            c := 0; extend version 2(ALL, 0, N)
          end output maximal complete subgraphs 2;

```

Remark on Algorithm 323 [G6]

Generation of Permutations in Lexicographic Order
[R.J. Ord-Smith, *Comm. ACM* 11 (Feb. 1968), 117]

Mohit Kumar Roy [Recd. 15 May 1972]
Computer Centre, Jadavpur University, Calcutta 32,
India

In presenting Algorithm 323, *BESTLEX*, for generating permutations in lexicographic order, the author has mentioned the number of transpositions. It may be remarked here that equal numbers of transpositions are required by both *BESTLEX* and the previously fastest algorithm, Algorithm 202 [1]. The exact number of transpositions (T_n) necessary to generate the complete set of $n!$ permutations is given by

$$T_n = n! (\psi_{n-1}) - (n+1)/2, \text{ if } n \text{ is odd, and}$$

$$T_n = n! (\psi_{n-2}) - n/2, \text{ if } n \text{ is even,}$$

$$\text{where } \psi_{2n} = 1 + \frac{1}{2!} + \frac{1}{4!} + \cdots + \frac{1}{(2n)!} \doteq 1.543 \text{ for } n \geq 3.$$

The above expressions do not include the few extra transpositions (equal to the integral part of $n/2$) required by *BESTLEX* to generate the initial arrangement from the final one, as this portion has

not been included in Algorithm 202. Therefore, the number of transpositions has no importance in the context of the claim that *BESTLEX* is more than twice as fast as Algorithm 202.

The main factor contributing to the speed of *BESTLEX* is the substantial reduction in the number of comparisons required, by the introduction of the own integer array q . Taking into account only those comparisons which involve array elements, the number of comparisons (C_n) required to generate all the $n!$ permutations can be shown to be equal to

$$C_n \text{ (Algorithm 202)} = \frac{n!}{2} [1 + 3\varphi_{n-2}] + n,$$

$$C_n \text{ (BESTLEX)} = n! [\frac{1}{2} + \varphi_{n-1}],$$

where $\varphi_n = 1 + \frac{1}{2!} + \frac{1}{3!} + \dots + \frac{1}{n!} \doteq 1.718$ for $n \geq 6$.

This shows that the number of comparisons required by *BESTLEX* is lower by .859($n!$) (approximately) in the case of the generation of all the $n!$ arrangements.

Finally, a modification of the *BESTLEX* algorithm is suggested which will reduce the number of comparisons again by $(n!)/2$. The modification involves replacement of lines 2-14 of Algorithm 323 by the following.

```
begin own integer array q[3:n]; integer k, m;
  real t; own Boolean flag;
comment Own dynamic arrays are not often implemented. The
  upper bound will have to be given explicitly;
if first then
begin first := false; flag := true
for m := 3 step 1 until n do q[m] := 1
end of initialization process;
if flag then
begin flag := false;
  t := x[1]; x[1] := x[2]; x[2] := t;
  go to finish
end;
flag := true;
for k := 3 step 1 until n do
```

References

1. Shen, Mok-Kong. Algorithm 202, generation of permutations in lexicographical order. *Comm. ACM* 6 (Sept. 1963), 517.

Added in proof: An improved version of *BESTLEX*, viz. Algorithm 323A, Generation of Permutation Sequences: Part 2, by R.J. Ord-Smith [*Comp. J.* 14, 2 (May 1971), 136-139], which also incorporates the modification suggested here, has come to the author's attention.

Remark on Algorithm 408 [F4]

A Sparse Matrix Package (Part I)

[John Michael McNamee, *Comm. ACM* 14 (Apr. 1971), 265-273]

E.E. Lawrence [Recd. 1 February 1972, 12 March 1973]
Central Application Laboratory, Mullard Limited,
New Road, Mitcham, Surrey CR4 4XY, England

The subroutines constituting Algorithm 408 were, with the exception of *MVSPMX* and *WRSPMX*, tested on an IBM 360/65 using CALL/360-OS. The author's alteration (iii) was introduced, i.e. declaration of the M -array to be half length. Other changes were introduced in order: (a) to make the algorithm more conversational in a time shared environment; and (b) to improve the speed of the sorting procedure in *PERCOL*.

The following deficiencies in the algorithm were noted

1. The dimensional parameters of *ACSPMX*, *ADSPMX*, and *MUSPMX* are incomplete. As an illustration of this consider the two matrices

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 4 \end{bmatrix}$$

$$B = \begin{bmatrix} 0 & 0 & 0 & 4 \\ 0 & 0 & 3 & 0 \\ 0 & 2 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

each of which has four nonzero elements.

Then the sum matrix has eight such elements, and in general, for two matrices with n_1 and n_2 nonzero elements, the number of nonzero elements, n_3 , in the sum matrix is in the range $0 \leq n_3 \leq n_1 + n_2$.

However in *ADSPMX* the condition used is $n_1 = n_2 = n_3$.

Similar arguments apply to *ACSPMX* and *MUSPMX*.

To correct this requires extensions to the parameter lists and dimension statements, and also it changes the conditional statements within the subroutines concerned.

This shows up with the CALL/360-OS system since the compiler performs subscript checking. It would not be evident on most compilers including the IBM Fortran IV G compiler. It is, however, bad practice to rely on default effects of compilers.

2. There are three, probably copying, errors in *MUSPMX* (page 270).

- (i) Line 33 should be:
IF(NCA.EQ.NCB) GO TO 3
- (ii) Line 55 should be:
DO 14 J = 1, NRB
- (iii) Line 102 should be:
CALL IPK(NRB,MC,2,NM)

Remark on Algorithm 420 [J6]

Hidden-Line Plotting Program

[Hugh Williamson, *Comm. ACM* 15 (Feb. 1972), 100-103]

Hugh Williamson [Recd. 9 Oct. 1972]

National Con-Serv, Incorporated, Austin, Texas

The input quantities to subroutine *HIDE* referred to in the following paragraphs (e.g. $N1$, *NFNS*, "input curve to be plotted") are described in the block of comment statements at the beginning of *HIDE* as originally published.

If $N1 < 0$, *DO* loop 71 is not executed properly, since the upper limit, $N1$, is less than the lower limit, 2. This affects only checking for monotonicity in the input abscissa array; otherwise, if the inputs are correct, the performance of the program is not affected.

The error is corrected if the first 11 executable statements are replaced by the following (the first executable statement of the original program, which is not changed, is listed for clarity):

```
IF(MAXDIM.LE.0) RETURN
IFPLOT = 1
IF(N1.GT.0) GO TO 76
N1 = -N1
IFPLOT = 0
```

Fig. 1. Without verticals.





Fig. 2. With verticals to aid visualization.



```

76 DO 71 I = 2,N1
   IF(X(I-1).LT.X(I)) GO TO 71
   MAXDIM = 0
   GO TO 75
71 CONTINUE
   IF(NG.GT.0) GO TO 5000

```

On computers in which all variables are not automatically set to zero before execution, *FNSM1* is not properly initialized if $NFNS \leq 0$. To correct this, simply insert the statement

$FNSM1 = 0$.

before the statement

```
IF(NFNS.LE.0) GO TO 46
```

The latter is the sixth statement after Fortran statement number 74.

FNSM1 will still be improperly defined if $NFNS = 1$. If only one curve is to be plotted, however, translating to simulate stepping in the depth dimension will not be done, so set $NFNS = -1$ for only one curve to be plotted.

In some cases, the three-dimensional surface is easier to visualize if (nearly) vertical lines are drawn at the left edge of each curve; this effect is illustrated by Figures 1 and 2. The verticals are added by inserting ($XMIN-\epsilon$, $YMIN$) as the first point in each input curve to be plotted, where ϵ is a small positive number ($10^{-4} \times DELTAX$ would be appropriate).

The author appreciates very much the comments received from readers of Communications regarding implementation of *HIDE* on different computers.

Remark on Algorithm 429 [C2]

Localization of the Roots of a Polynomial [C2]
[W. Squire, *Comm. ACM* 15 (Aug. 1972), 776-777]

H.B. Driessen and E.W. LeM. Hunt [Recd. 13 Oct. 1972, 29 Jan. 1973]

Supreme Headquarters Allied Powers of Europe,
Technical Center, P.O. Box 174, The Hague, The
Netherlands

There seems to be an error in this algorithm. If we take the polynomial:

$$z^4 + a_2 z^2 + a_3 z^3 + a_4 z + a_5 = 0,$$

then after the second pass through the K -loop of the logical function *HRWTZR*(C, N), the term $(a_2 a_3 - a_4) a_4 - a_5 a_2$ is tested for a minus sign. However, the term which should be tested according to the Routh-Hurwitz criterion is $(a_2 a_3 - a_4) a_4 - a_5 a_2^2$. If this term is negative then there are no roots with positive real parts.

As an example, if the polynomial

$$z^4 + 5.6562 z^3 + 5.8854 z^2 + 7.3646 z + 6.1354 = 0$$

is studied with the help of Algorithm 429 one will find as output:

Roots are in an annulus of inner radius $.454 E + 00$ and outer radius $.836 E + 01$;

There are no real positive roots;

The negative roots (if any) are between $-.454 E + 00$ and $-.836 E + 01$;

There are no roots with positive real parts.

However, if one calculates the roots of this equation, one will find approximately:

$$\begin{aligned} z_1 &= -1.0001 \\ z_2 &= -4.7741 \\ z_{3,4} &= +0.0089 \pm 1.1457 i \end{aligned}$$

Statement 20 + 1 in the logical function *HRWTZR*(C, N), which was originally " $C1 = C(1)$ ", should be amended to read " $C1 = C(1)/C1$ ".

As a by-product of our investigation, it turns out that the structure of the logical function *HRWTZR* can be simplified by abandoning the logically redundant steps $C(K) = C(K+1)$.

The following listing incorporates both the correction and the simplifications. The function has been parameter tested on a CDC-6400.

```

LOGICAL FUNCTION HRWTZR (C,N)
DIMENSION C(N)
HRWTZR = .FALSE.
IF (C(1) .LE.0 .OR. C(N) .LE.0) RETURN
C1 = C(1)
M = N - 1
DO 30 I = 2,M
DO 20 K = I,M,2
20 C(K) = C(K) - C(K+1)/C1
C1 = C(1)/C1
IF (C1.LE.0) RETURN
30 CONTINUE
HRWTZR = .TRUE.
RETURN
END

```