Appl. Statist. (1977),
26, *No.* 3, *pp.* 251–260

A General Algorithm for Analysis of Variance

By R. W. PAYNE and G. N. WILKINSON†

Rothamsted Experimental Station, Harpenden, Herts, Britain

[Received August 1976. Revised March 1977]

SUMMARY

The general method of analysis of variance described by Wilkinson (1970) is reviewed. The implementation of an algorithm based on this method is described with particular attention to the doubly linked binary tree used to store information about the design.

Keywords: ANALYSIS OF VARIANCE; EFFICIENCY FACTOR; ANALYSIS OF VARIANCE ALGORITHM; COMPUTER PROGRAMMING

1. INTRODUCTION

WILKINSON (1970) described a general analysis-of-variance method which enables a single algorithm to analyse most common experimental designs. These designs include all orthogonal designs (possibly with several error terms), designs with balanced confounding and other designs with a specified pattern of confounding. The normal equations are not solved directly but by sequentially "sweeping" a working variate which initially contains the data values. The sweep operation calculates effects for a particular term of the linear model from the working variate and uses them to modify it as described below. The algorithm has been implemented as part of the Genstat system (Wilkinson and Rogers, 1974; Nelder *et al.*, 1975). In this paper we review the methods of analysis, explaining how the necessary sequence of sweeps is determined automatically, so that the user only defines the error and treatment terms to be included (together with the treatment combinations applied to each experimental unit). We describe some of the facilities of the program, and how they are implemented, paying particular attention to how information about the design is stored.

2. DESCRIPTION OF THE ALGORITHM

2.1. *The Method of Analysis*

The algorithm fits linear factorial models like

$$y_{ijk} = \mu + (\delta_i + \omega_{ij} + \varepsilon_{ijk}) + (\alpha_j + \beta_k + \gamma_{jk}). \tag{2.1}$$

This represents the model for a split-plot design with *block* (i.e. error) *terms* δ_i, ω_{ij}, ε_{ijk} (in the first bracket) representing block residuals, whole-plot residuals and sub-plot residuals respectively, while the *treatment terms* α_j, β_k and γ_{jk} (in the second bracket) represent the main effects of factors A and B, and their interaction. The model is specified as described by Wilkinson and Rogers (1973) and interpreted by an extended version of algorithms AS 65, Rogers (1973), with block terms specified separately from the treatment terms. The grand mean, μ, which is part of both block and treatment formulae, is included in the model automatically.

The model is fitted incrementally and the only constraints imposed on terms in it are those generated by earlier fitted terms. Thus terms are deleted that are marginal to preceding terms in a formula, e.g. $A.B+A$ becomes $A.B$. The model formula for the treatment terms of (2.1) is $A+B+A.B$ so that the main effects A and B are fitted before $A.B$ which thus represents the

† Present address: Mathematics Research Centre, Madison, Wisconsin, U.S.A.

interaction effects γ_{jk}. If B were deleted from the model, $A.B$ would represent B-within-A effects specifying a nested factorial model.

It is well known that orthogonal models with only one error term can be fitted by *sweeping* a working variate (which initially contains the data values) for each treatment model term, including μ. By *sweep* is meant a two-part operation on the values of the working variate, as specified by a model term. In the first part, termed a *forward sweep*, the means are calculated for every combination of the levels of the factors in the term, and used to calculate effects for the term. For orthogonal designs the effects are the actual means but for non-orthogonal designs the means are divided by an *efficiency factor* (see below). In the second part, termed a *backward sweep*, the effects are subtracted from the working variate.

As an example consider the sweep for B in

$$y_{ijk} = \mu + (\varepsilon_{ijk}) + (\alpha_j + \beta_k + \gamma_{jk}). \tag{2.2}$$

If, after the sweeps for μ and A, the working variate contains values y'_{ijk}, the first stage of the sweep estimates each β_k by $y'_{..k}$; subtracting these estimated effects, y'_{ijk} then becomes $y'_{ijk} - y'_{..k}$. At each stage the variate contains the residual effects after fitting the preceding terms; thus it finally contains the estimated residuals ε_{ijk}.

A modification enables models with more than one error term to be fitted. For (2.1), the analysis proceeds as before, initially sweeping out μ, δ_i and ω_{ij} to leave only subplot effects. These are partitioned into treatment and residual effects by performing a sweep for each treatment model term. As A has been applied to whole plots, there is no point in making a sweep for A at this stage. (We describe below how redundant sweeps are detected.) Thus sweeps for B and $A.B$ give the within-whole-plot analysis. The between-whole-plot analysis is obtained by replacing each value in the working variate by the corresponding whole-plot effect calculated in the initial sweep for whole-plots. Thus if the initial sweep produced estimated effects ω'_{ij}, these are placed in the ijkth elements of the working variate. This operation is termed a *pivot* for whole plots. Whole-plot effects were calculated before any treatment effects were removed; they thus contain the effects of any treatment terms estimated between whole plots as well as the whole-plot residuals (but not block effects which had already been removed). A sweep for A thus completes the between-whole-plot analysis leaving the whole-plot residuals. There are no treatment effects to be estimated between blocks so the analysis is complete.

The total sum of squares has thus been partitioned into three components, referred to as *strata*, one for each block term. The total sum of squares for each stratum is subdivided into sums of squares for treatment terms estimated in that stratum, and the stratum residual. This method of analysis ensures that mean squares for treatment terms are automatically compared with the correct error mean square (the appropriate stratum residual).

We now show how models containing non-orthogonal terms can be analysed. Wilkinson (1970) uses matrix notation to describe the method. Here we give a geometrical description which may provide more insight. Suppose A and B in 2.2 are non-orthogonal and that each has only one degree of freedom. The data space is defined to be the n-dimensional space (where n is the number of units y_{ijk}) representing all possible values of the data vector $\mathbf{y}$. If levels 1 and 2 of A have replication r_1 and r_2 respectively then α_1 and α_2 are of the form $r_2 a$, $-r_1 a$ for some real number a; i.e. the effect of A can be represented in the data space by a vector $a\mathbf{d}_A$: where $(\mathbf{d}_A)_m$ is r_2 if A takes level 1 on the mth experimental unit, and $-r_1$ if A takes level 2. Thus, for A with one degree of freedom, the subspace of possible A effects is a line. B effects are similarly represented by vectors $b\mathbf{d}_B$. Fig. 1 shows the plane spanned by vectors $\mathbf{d}_A, \mathbf{d}_B$. $\mathbf{d}_A^{\perp}$ and $\mathbf{d}_B^{\perp}$ are vectors in the plane orthogonal to $\mathbf{d}_A$ and $\mathbf{d}_B$ respectively. Vector $\mathbf{y}'$ is the projection of the data vector onto the plane. If A is fitted first, the appropriate sweep projects $\mathbf{y}'$ onto $\mathbf{d}_A^{\perp}$ by subtracting the vector **31**—the effect of A ignoring B. Thus $\mathbf{y}'$ is transformed from vector **01** to vector **03**. A similar sweep for B would estimate the B effect as

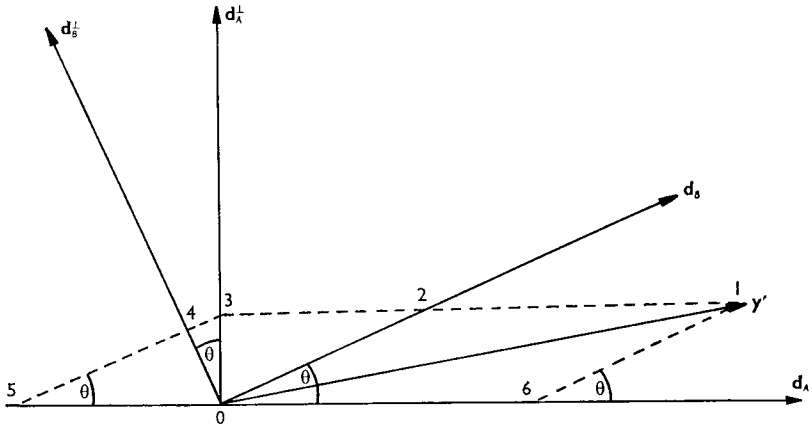

FIG. 1. A geometric representation of the sweeps required to analyse non-orthogonal terms A and B.

the vector **43** and would further transform $\mathbf{y}'$ to **04** on $\mathbf{d}_B^\perp$. However, the effect of B eliminating A is **53** so the projection should be continued onto $\mathbf{d}_A$, transforming $\mathbf{y}'$ to **05**. Thus the B effect should be estimated as an *effective mean* **53** (not as the simple mean **43**); this is the simple mean divided by e, where $e = \|\mathbf{43}\|/\|\mathbf{53}\| = \sin\theta x\|\mathbf{03}\|/\|\mathbf{53}\| = \sin^2\theta$; e is the efficiency factor defined by Yates (1936). If the two factors are orthogonal, $\theta = \pi/2$, i.e. $e = 1$; whereas if they are aliased $\theta = 0$, i.e. $e = 0$. The analysis is completed by a *reanalysis sweep* for A which, projecting once more onto $\mathbf{d}_A^\perp$, subtracts the vector **05** leaving a null vector. (The effect of A eliminating B is $\mathbf{06} = \mathbf{31} + \mathbf{05}$.) The sum of squares for A ignoring B is given by $\|\mathbf{31}\|^2$, the sum of squares of the A effects; but the sum of squares for B eliminating A is given by $\|\mathbf{03}\|^2$, the sum of squares of the B effects, multiplied by the efficiency factor.

This method of analysis can also be used if A or B has more than one degree of freedom, provided all the B contrasts have the same efficiency factor, i.e. provided the angle between the A subspace and each of the vectors spanning the B space is θ. Such terms are said to be *balanced*. In designs with more than one model term, the derivation of the efficiency factor is similar to that above, except that the subspace generated by A effects becomes the subspace generated by all effects preceding B in the model formula. If B is non-orthogonal to more than one preceding term, a reanalysis sweep is required for each such term. If any one of the reanalysis terms is itself non-orthogonal to preceding terms, reanalysis sweeps will be required after the sweep for that term, and so on. In practice the algorithm does not allow sequences of reanalysis sweeps to be nested more than ten deep but this limit should not be exceeded in a designed experiment. If a treatment term is non-orthogonal to a block term, the reanalysis sweep in the stratum for that block term consists of a recalculation of the effective block means and then a pivot to reset each of the observations in the working variate to the corresponding mean. This is termed a *forward pivot*.

The restriction that all effects of a non-orthogonal factor must have the same efficiency factor can be made less stringent by defining pseudo-terms linked to a treatment model term. In the analysis each pseudo-term is examined separately but the sums of squares are recombined when forming tables of means. Thus pseudo-terms should be used to group contrasts estimated in the same stratum with the same efficiency factor, i.e. to represent subspaces each of which makes a single angle with the subspace of preceding terms.

2.2. *Determining the Structure of a Design*

We now explain how efficiency factors can be calculated and describe the *dummy analysis* used to determine the design structure.

Consider non-orthogonal factors A and B, as before, each with only one degree of freedom, and any vector $\mathbf{y}$ which contains non-zero effects for both A and B. Fig. 2 again shows the data space projected onto the subspace of A and B effects and $\mathbf{y}'$ is the corresponding projection of $\mathbf{y}$. If we sweep for A and B as though the design were orthogonal, $\mathbf{y}'$ is transformed from

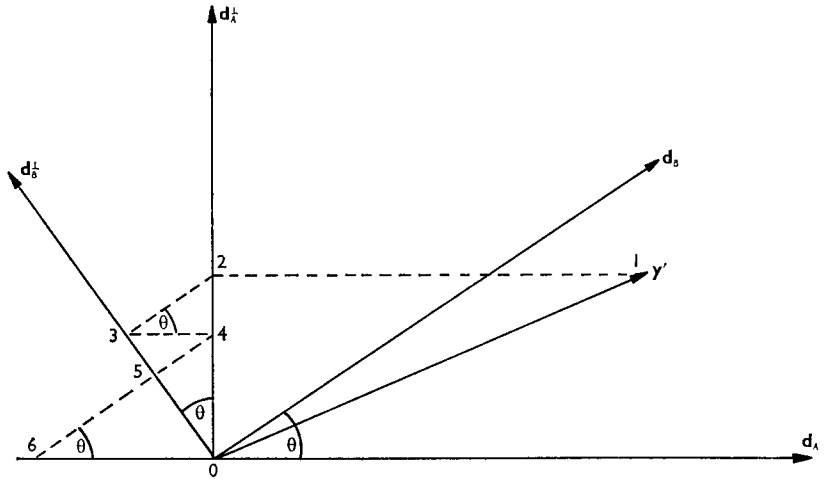

FIG. 2. A geometric representation of the sweeps required to determine the efficiency factor of B.

01 to **03**. If the process is repeated, $\mathbf{y}'$ is further transformed to **05**. The efficiency factor can now be determined from the originally calculated B effect and the B effect in the second sweep, since

$$\|\mathbf{54}\| = \|\mathbf{34}\| \times \cos\theta = \|\mathbf{32}\| \times \cos^2\theta \Rightarrow e = 1{\cdot}0 - \|\mathbf{54}\|/\|\mathbf{32}\|.$$

If A and B are orthogonal, the second sweeps for A and B produce zero sums of squares; and if B has more than one degree of freedom, each effect has the same efficiency factor, if and only if the ratio of each effect in the first sweep to the corresponding effect in the second sweep is $\|\mathbf{32}\|/\|\mathbf{54}\|$. Thus both orthogonality and balance can be checked. If in the second sweep for B, the effects are merely calculated but not subtracted from the working variate (i.e. a forward sweep), the efficiency factor can be calculated and the effects adjusted before being subtracted from the variate (by a backward sweep). The vector $\mathbf{y}'$ is then transformed to **06**; thus a further (reanalysis) sweep for A completes the analysis for A and B. This method, used in the present dummy analysis, is shown by Rogers and Wilkinson (1974) to have better numerical properties than the original method of Wilkinson (1970).

Since any vector $\mathbf{y}$ can be used provided it does not have a zero sum of squares for any model term, the same variate can be used for the whole dummy analysis with each model term examined in turn. The dummy variate used at present contains uniform random numbers and has proved satisfactory with the designs so far analysed. The analysis starts with a sweep for the grand mean and then sweeps to examine first the block terms and then the treatment terms. If a model term is not estimable, i.e. is aliased with preceding terms, its sum of squares is zero apart from numerical round-off, and the next term is examined. When a term is found with a non-zero sum of squares, the algorithm reanalyses for all previous terms, noting any with a non-zero sum of squares as non-orthogonal to the new term, and then uses a forward sweep for the term to determine the efficiency factor. If the new term is orthogonal to the previous terms, the algorithm can proceed to the next term; otherwise a backward sweep for the term is needed first, followed by reanalysis sweeps for the terms to which it is non-orthogonal.

This process determines which treatment terms are estimable in the bottom stratum and calculates their efficiency factors. The algorithm then works back through the block formula examining the stratum generated by each term as follows. Each stratum analysis starts with a pivot and then the treatment terms are examined as before, except that the reanalysis sweep for the stratum term is a forward pivot instead of an ordinary sweep and that there need be no reanalysis sweeps for lower strata.

Thus the dummy analysis is very like a data analysis except that all possible sweeps are attempted and extra sweeps are necessary to determine the efficiency factors.

2.3. *Additional Facilities*

The method described above can easily be extended to allow weighted analysis of variance by estimating the effects from effective weighted means (see Wilkinson, 1970).

Treatment terms can be partitioned into regression components. The x-variates can be either provided explicitly or generated automatically for polynomial regression. To simplify the analysis, the x-variates for each treatment term are orthogonalized so that the fit of the second x-variate represents the deviation of the fit of the first two x-variates from the fit of the first variate, and so on. Thus the x-variates for polynomial regression are orthogonal polynomials. The regression calculations use operations similar to the sweeps described above, operating on a y-variate of estimated treatment effects. Each sub-model sweep calculates the regression coefficient of the x-variate concerned and then subtracts the regression effects from the y-variate; thus, after the regression sweeps, the y-variate contains deviations from the regressions. (The regressions should be weighted according to the replication of the effect concerned.)

Covariance analysis is implemented separately using the residuals, effects, efficiency factors, etc. produced by the algorithm.

3. Implementation of the Algorithm

The algorithm has four phases which (i) interpret the model, (ii) interpret the design, (iii) analyse the data variates and any covariates and (iv) print the results.

Each of the first three phases can be regarded as a process of obtaining information and storing it for future use. With an algorithm as general in scope as this, some form of dynamic storage allocation is essential. This is achieved by setting up various structures in a workspace array. Each structure contains information about one aspect of the model. Most of the structures are fairly simple in form. For example, the structure formed when the factors in the design are checked contains lists of the number of levels of each factor, lists of origins of factor values, etc. Information about each factor is thus easily accessed using the index of the factor and the list origins.

3.1. *The Design Control Structure*

In the design control structure, there is no obvious implicit indexing since a term can be estimated in more than one stratum. To avoid storing extra indexing information, information about analysis terms is stored in the nodes of a doubly linked binary tree, formed during the dummy analysis (see Section 3.2). A binary tree is a hierarchical structure in which each node may point to at most two other nodes. In a doubly linked tree each node also points back to the previous node (see Knuth, 1969). The first node in the tree, known as the root, stores information about the grand mean.

The first branch from the root points to a node for the first block term and the second branch is empty. The first branch from each block node points to a node for the next block term (if any) and the second branch points to the first of a list of nodes for treatment terms estimated in that stratum, in the order in which they are estimated. If no treatment terms are estimated in the stratum, the second branch will be empty. Nodes for pseudo-terms of a

model term are preceded by what is known as a *group term* node. The first branch from that node points to the next treatment term and the second branch points to the list of pseudo terms. This serves to indicate that the pseudo-terms are all linked to one model term. The first branch from a treatment term or a pseudo-term points to the next treatment term or pseudo-term (if any), and the second branch is empty. Treatment terms thus have a node for each stratum in which they are estimated. Each node contains the efficiency factor; a balance indicator (0 for an orthogonal term, 1 for a non-orthogonal term and -1 for a group term); addresses of the previous node and the two subsequent nodes; indices and addresses for accessing information about the term in the other structures; and a list of addresses of nodes for previous terms non-orthogonal to the term concerned (i.e. those for which reanalysis sweeps are required).

This is better understood by considering an example. Table 1 shows the design, analysis-of-variance table and efficiency factors for a factorial design with two factors—*A* with two levels and *B* with four levels—in eight blocks of four plots. The Block formula is thus

BLOCKS + BLOCKS.PLOTS

TABLE 1

The design, analysis of variance table and efficiency factors for a factorial design with factors A and B in eight blocks of four plots. Each degree of freedom of A.B is confounded in two of the eight blocks. A and two degrees of freedom of B are completely estimated in the Blocks.Plots stratum, whereas the other degree of freedom of B—represented by PF—is estimated in both strata

	Plot 1			*Plot* 2			*Plot* 3			*Plot* 4		
	A	*B*	*PF*	*A*	*B*	*PF*	*A*	*B*	*PF*	*A*	*B*	*PF*
Block 1	1	1	2	2	2	1	2	3	1	1	4	2
Block 2	2	1	2	1	2	1	1	3	1	2	4	2
Block 3	2	1	2	1	2	1	2	3	1	1	4	2
Block 4	1	1	2	1	3	1	2	2	1	2	4	2
Block 5	2	1	2	1	3	1	2	2	1	1	4	2
Block 6	1	1	2	1	2	1	2	3	1	2	4	2
Block 7	1	2	1	1	3	1	2	2	1	2	3	1
Block 8	1	1	2	2	1	2	1	4	2	2	4	2

Analysis of variance		*Efficiency factors*	
Source of variation	D.F.	Model term	E.F. non-orthogonal terms
Blocks stratum		Blocks stratum	
B	1	*PF*	0·250
A.B	3	*A.PF*	0·250
Residual	3	*A.B*	0·250
Total	7		
Blocks.Plots stratum		Blocks.Plots stratum	
A	1	*A*	1·0
B	3	*PF*	0·750 Blocks
A.B	3	*B′*	1·0
Residual	17	*A.PF*	0·750 Blocks
Total	24	*A.B′*	0·750 Blocks
Grand total	31		
Total number of observations	32		

Each degree of freedom of *A.B* is confounded in two of the eight blocks, thus *A.B* has efficiency factor 0.25 in the Blocks stratum and 0·75 in the Blocks.Plots stratum. The main effect of *A* and two degrees of freedom of *B* are completely estimated in the Blocks.Plots stratum while

the other degree of freedom of B is estimated in both strata and must thus be specified by the pseudo-term, PF. The treatment formula is

$$A+PF+B'+A.PF+A.B',$$

where B' and $A.B'$ represent all effects of B and $A.B$ not already specified by the preceding pseudo-terms. Table 2 shows the design control structure for the analysis (omitting irrelevant

TABLE 2

The design control structure for the design in Table 1

Node	*Model term*	*Efficiency factor*	*Balance indicator*	*Previous node*	*Next node* (1)	*Next node* (2)	*Nodes of non-orthogonal terms for reanalysis sweeps*
1	Grand mean	1·0	0 (orthogonal)	0	2	0	
2	Blocks	1·0	0	1	3	11	
3	Blocks .Plots	1·0	0	2	0	4	
4	*A*	1·0	0	3	5	0	
5	Group term for *B*	—	−1 (group term)	4	8	6	
6	*PF*	0·75	1 (non-orthogonal)	5	7	0	2 (Blocks)
7	*B'*	1·0	0	6	0	0	
8	Group term for *A.B*	—	−1	5	0	9	
9	*A.PF*	0·75	1	8	10	0	2 (Blocks)
10	*A.B'*	0·75	1	9	0	0	2 (Blocks)
11	Group term for *B*	—	−1	2	13	12	
12	*PF*	0·25	1	11	0	0	2 (Blocks)
13	Group term for *A.B*	—	−1	11	0	14	
14	*A.PF*	0·25	1	13	15	0	2 (Blocks)
15	*A.B'*	0·25	1	14	0	0	2 (Blocks)

accessing information for other structures). Fig. 3 shows the linkage between nodes in diagrammatic form. There are two block terms, Blocks (node 2) and Blocks.Plots (node 3). The first branch from node 2 points to node 3, whereas as Blocks.Plots is the bottom stratum, the first branch from node 3 is empty. The second branch from node 3 points to node 4 for A, the first treatment term estimated in the bottom stratum. The second branch from node 2 points to node 11, the group term for B. Branch one from the group term node points to the next treatment term in the stratum; branch two points to a node for PF, the pseudo-term of B, estimated in the Blocks stratum.

With the nodes linked in this way, the form of the tree represents the analysis sequence. The accessing subroutines allow nodes to be accessed sequentially in one of two orders. The order of nodes used when the data are analysed starts at the root and then takes the first branch at every branching node until, on reaching the last block term, the first branch is empty. It then works back up the tree and takes the second branch at each block term node to examine

the treatment terms estimated in the stratum concerned, and finally terminates at the root. The accessing routine indicates whether a block term has been encountered on the way from the root, when a standard sweep is required, or on the way back to the root, when a pivot is required before the treatment terms in the stratum are analysed. If a group term is encountered

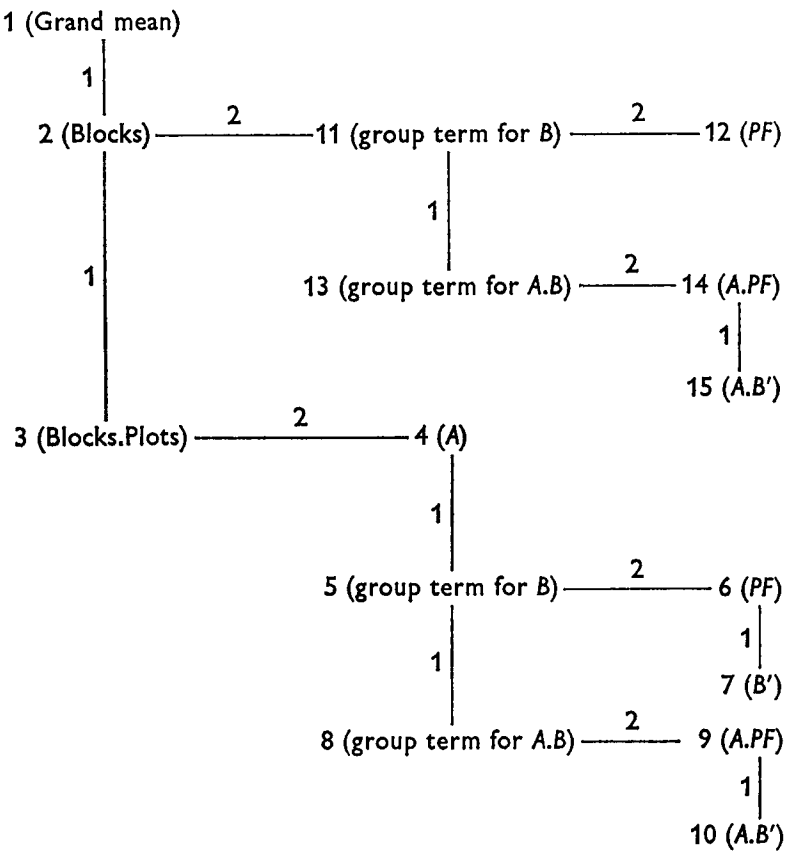

FIG. 3.

in a list of treatment terms, the pseudo-terms (on the second branch) are examined before subsequent treatment terms. If a treatment term is non-orthogonal to preceding terms, nodes for these terms are recovered in turn with a parameter set to indicate a reanalysis sweep, and each reanalysis term is followed by its own reanalysis sequence if necessary. Thus the order of nodes for the example above would be:

1 (Grand mean sweep);
2 (sweep for Blocks);
3 (initialize for the analysis of the Blocks.Plots stratum);
4 (sweep for *A*);
5 (group term for *B*—no action);
6 (sweep for *PF*);
2 (reanalysis sweep for Blocks);
7 (sweep for *B'*);
5 (return to group term for *B*—no action);

8 (group term for $A.B$—no action);
9 (sweep for $A.PF$);
2 (reanalysis sweep for Blocks);
10 (sweep for $A.B'$);
2 (reanalysis sweep for Blocks);
8 (return to group term—no action);
3 (store Blocks.Plots residuals);
2 (forward pivot for Blocks);

followed by the analysis for the Blocks stratum. This analysis is similar to that for the Blocks.Plots stratum except that there is no sweep for B' and the reanalysis sweeps for Blocks take the form of forward pivots;

2 (store Blocks stratum residuals);
1 (end of analysis).

A different order of scanning is used in printing the analysis-of-variance table, when the breakdown of the sum of squares is given for each stratum in turn, in the order in which the terms occur in the block formula. Thus, after each node for a block term, the second branch is taken first, to access the treatment terms estimated in the associated stratum; then the first branch is taken and the stratum below is considered. The order of branching at the group term is as before. For the example this would be:

1 (Grand mean—print heading);
2 (print Blocks stratum heading);
11 (group term for B—set accumulators for the sum of squares and degrees of freedom for B);
12 (add S.S. and D.F. for PF to accumulators);
11 (return to group term for B—print line for B in the analysis of variance table);
13, 14 and 15 calculate and print the corresponding line for $A.B$;
2 (print residual and total lines for Blocks stratum);

followed similarly by the nodes for terms in the Blocks.Plots stratum. This illustrates why nodes for group terms are necessary. This order is also used for printing the tables of effects and residuals.

3.2. *Forming the Design Control Structure*

The structure is formed sequentially during the dummy analysis (2.2), starting with the root, set up after the initial grand-mean sweep. After a term has been detected, the sweeps required to reanalyse previous terms are determined by scanning the partially completed tree in a similar order to that used in the data analysis; i.e. starting at the root and taking the first branch at each block node until, after reaching the node for the stratum concerned, the second branch is taken to examine previous treatment terms in that stratum. Thus, in the example above, the order of nodes after term PF has been detected in the Blocks.Plots stratum would be:

1 (Grand mean—this sweep is not strictly necessary but helps to prevent numerical round-off building up in the working variate);
2 (sweep for Blocks);
3 (Blocks.Plots, no action since pivots are not required for the bottom stratum);
4 (sweep for A);

followed by a forward sweep for PF. The sweep for Blocks and the sweep for PF both produce non-zero sums of squares, thus PF is discovered to be non-orthogonal to Blocks and the efficiency factor is calculated see (2.2). A node for PF can now be added at the end of the branch; however, since PF is a pseudo-term, node 5—the group term for B—is added first. Then, after a backward sweep for PF, the next term—B'—is examined. When all the terms have been examined, the scan moves back to the previous block node to determine the treatment terms estimated in that stratum.

4. Conclusion

The methods described above allow a great many designs to be analysed by a single program. Such a program will not, of course, be as efficient in computer time or storage for any given design as a program written especially for that design or restricted class of designs. However, unless the same design is to be used very many times, such savings in computer costs will be small compared to the programming cost. The algorithm, as implemented in Genstat, has been used successfully for routine analysis of variance at Rothamsted and other A.R.S. institutes for the last 6 years. In GENSTAT a small analysis on the Rothamsted ICL 4-70 requires approximately 100 K bytes of computer storage.

Acknowledgement

We acknowledge very helpful advice on presentation from D. A. Preece.

References

KNUTH, D. E. (1969). *The Art of Computer Programming*, Vol. 1, pp. 228–463. Reading, Mass.: Addison-Wesley.

NELDER, J. A. *et al.* (1975). Genstat Reference Manual. Inter University/Research Council Series Report No. 3, 3rd edn.

ROGERS, C. E. (1973). Algorithm AS 65. Interpreting structure formulae. *Appl. Statist.*, **22**, 414–424.

ROGERS, C. E. and WILKINSON, G. N. (1974). Numerical problems in recursive analysis of variance algorithms. *Bull. Inst. Math. and Applicns*, **10**, 141–143.

WILKINSON, G. N. (1970). A general recursive algorithm for analysis of variance. *Biometrika*, **57**, 19–46.

WILKINSON, G. N. and ROGERS, C. E. (1973). Symbolic description of factorial models for analysis of variance. *Appl. Statist.*, **22**, 392–399.

—— (1974). Genstat Users' Guide No. 2. Analysis of Designed Experiments, Inter University/Research Council Series Report No. 16.

YATES, F. (1936). Incomplete randomised blocks. *Ann. Eugen.*, **7**, 121–140.