31

Reducing the Number of AD Passes for Computing a Sparse Jacobian Matrix

Shahadat Hossain and Trond Steihaug

ABSTRACT A reduction in the computational work is possible if we do not require that the nonzeros of a Jacobian matrix be determined directly. If a column or row partition is available, the proposed substitution technique can be used to reduce the number of groups in the partition further. In this chapter, we present a substitution method to determine the structure of sparse Jacobian matrices efficiently using forward, reverse, or a combination of forward and reverse modes of AD. Specifically, if it is true that the difference between the maximum number of nonzeros in a column or row and the number of groups in the corresponding partition is large, then the proposed method can save many AD passes. This assertion is supported by numerical examples.

31.1 Introduction

To determine a sparse Jacobian matrix, one can compute a partition of the columns and then use finite differences (FD) or automatic differentiation (AD) to obtain the nonzeros. In a direct method [1, 4, 6, 11], the nonzero entries corresponding to a group of columns can be read off the difference approximations or a forward pass directly, i.e., without any other arithmetic operations. In an indirect method such as a substitution method the unknowns are ordered such that the ordering leads to a triangular system [2, 4, 10]. The additional cost of substitution is often small compared with the savings in the number of AD passes. Substitution methods have been investigated mainly in the context of symmetric matrices [2, 9, 14]. In this chapter we extend the results in [10] and present numerical experiments. The techniques presented here assume that a column (row) partition or a complete direct cover [11] or bipartition [5] is available. Using the sparsity information we define systems of linear equations that can be solved by substitution. If the number of groups in the partition is larger than the maximum number of nonzeros in any row of the Jacobian matrix, then the proposed techniques reduce the number of AD passes by at least one.

In this chapter, §31.2 begins with a brief discussion of the CPR method

[6]. We show how to reduce the number of AD passes by defining a new grouping of columns from a given column partition. Our main result is presented in this section. We illustrate the usefulness of the proposed techniques by a motivating example. In §31.3, we show that under certain sparsity restrictions, the number of AD passes can be reduced further by grouping together columns from more than two structurally orthogonal column groups into one. §31.4 presents the numerical experiments. In the final section, we review the methods presented in the chapter and discuss future research directions.

In this chapter if a capital letter is used to denote a matrix (e.g., A, B, X), then the corresponding lower case letter with subscript ij refers to the (i, j) entry (e.g., a_{ij}, b_{ij}, x_{ij}). The elements in an arbitrary row will be denoted by Greek letters.

31.2 Partitioning and Substitution Ordering

Assume that we have a priori knowledge of the sparsity structure of the Jacobian matrix $A \in R^{m \times n}$. We want to obtain a *seed matrix* $X \in R^{n \times p}$ such that nonzeros of A can be recovered by substitution with p as small as possible. Solve for nonzero a_{ij} in

$$X^T A^T = B^T, \tag{31.1}$$

where the columns of B are matrix-vector products of A computed from one AD pass or a difference in function evaluations $B \leftarrow AX$. A group of columns are said to be *structurally orthogonal* if they do not contain nonzeros in the same row position. In the CPR approach columns of A are partitioned into structurally orthogonal groups. The resulting partition is called a *consistent partition.* Corresponding to each group of structurally orthogonal columns $C_k,\ k = 1, 2, \ldots, p$, where C_k refers to a set of column indices, we define $x_{jk} = 1$ if column j of A is in C_k and $x_{jk} = 0$ otherwise. Then, (31.1) defines m diagonal linear systems in the nonzeros of each row of A.

If we do not require that the nonzeros be determined directly, it may be possible to reduce the computational work further. To illustrate, we consider the following matrix

$$A = \begin{bmatrix} a_{11} & 0 & a_{13} \\ a_{21} & a_{22} & 0 \\ 0 & a_{32} & a_{33} \end{bmatrix}, \text{ and choose } X = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

For example, the second row of A can be determined by solving for a_{21} and a_{22}

$$\begin{bmatrix} a_{21} & a_{22} & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} b_{21} & b_{22} \end{bmatrix}.$$

Eliminating row 3 of X and transposing the system, we get

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a_{21} \\ a_{22} \end{bmatrix} = \begin{bmatrix} b_{21} \\ b_{22} \end{bmatrix},$$

an upper triangular system. The nonzeros in two other rows of A can be solved for in an analogous way, requiring only two AD passes. On the other hand, the associated column intersection graph $G(A)$ is a complete graph, and hence the columns have only a trivial partition of 3 structurally orthogonal groups requiring 3 AD passes.

Assume that $A \in R^{m\times n}$ with at most ρ nonzeros per row. Let the columns be partitioned into $p \geq \rho + 1$ groups C_i, $i = 1, 2, \ldots, p$. Consider a new grouping of the columns,

$$\mathcal{C}_i = \{C_i, C_{i+1}\}, \quad i = 1, 2, \ldots, p - 1. \tag{31.2}$$

In the above, we have adopted the notation $\{C_i, \ldots, C_k\} = C_i \cup \cdots \cup C_k$. We claim that (31.2) defines a substitution ordering of the unknowns. Without loss of generality, let $p = \rho + 1$, and let $\alpha_1, \alpha_2, \ldots, \alpha_p$ denote the elements in row i of A that we want to determine. Then

$$\begin{bmatrix} 1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 1 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & 0 \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 1 & 1 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \vdots \\ \alpha_{p-1} \\ \alpha_p \end{bmatrix} = \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \vdots \\ \beta_{p-1} \end{bmatrix}, \tag{31.3}$$

where $\beta_1, \ldots, \beta_{p-1}$ is row i in B (31.1). Since there are at most $\rho < p$ nonzeros in any row of A, one of the variables α_k, $k \in \{1, 2, \ldots, p\}$ in (31.3) is zero, so that at least one of α_{k-1} and α_{k+1} can be determined directly, and the remaining nonzeros can be determined by substitution. The above discussion establishes the following result.

Lemma 1 *If $p > \rho$ then $\{C_1, C_2\}, \{C_2, C_3\}, \cdots, \{C_{p-1}, C_p\}$ defines a substitution ordering of the nonzero elements of A.*

An idea along the lines of Lemma 1 has been considered in [13] as "multi-coloring". Lemma 1 shows that the substitution ordering saves one AD pass provided that $p > \rho$. This result also indicates that the problems where the Jacobian matrix has a relatively few nonzeros per row but the chromatic number of its intersection graph is large [15] are good candidates for a substitution scheme.

If each row of the compressed matrix has at least two consecutive zero elements, then the matrix formed by $\{C_1, C_2\}, \{C_2, C_3\}, \cdots, \{C_{p-1}, C_p\}$ will have at least one zero element in every row, and the process could be repeated. However, as will be shown in the next section, merging of the successive index sets is done only once.

We note that any arbitrary combination of $\rho + 1$ column groups saves one AD pass. If $p \; modulo \; (\rho+1) \neq 0$, then it is likely that some ρ or fewer column groups has to be determined directly. Thus, the total number of AD passes in the substitution method is at most $p - \lfloor \frac{p}{(\rho+1)} \rfloor$.

Let us assume that a consistent partition of the columns of A contains

$$p = k(\rho + 1)$$

groups, where $k \geq 1$. Then we save k forward passes.

We illustrate the above substitution method on an example. Consider a 15×6 matrix A with the following sparsity pattern [15].

$$\begin{bmatrix} \times & \times & & & & \\ & & \times & \times & & \\ & & & & \times & \times \\ \times & & \times & & & \\ \times & & & \times & & \\ & \times & \times & & & \\ & \times & & \times & & \\ \times & & & & \times & \\ \times & & & & & \times \\ & \times & & & \times & \\ & \times & & & & \times \\ & & \times & & \times & \\ & & \times & & & \times \\ & & & \times & \times & \\ & & & \times & & \times \end{bmatrix} \tag{31.4}$$

The column intersection graph $G(A)$ is a graph on 6 vertices, and the chromatic number is 6, hence any consistent column partition must have at least 6 column groups. Since $\rho = 2$ for this example we can define the two parts

$$\{C_1, C_2\}, \; \{C_2, C_3\} \text{ and } \{C_4, C_5\}, \; \{C_5, C_6\}, \tag{31.5}$$

where C_j contains column j of A for $j = 1, 2, \ldots, 6$. Using our substitution method, we need only 4 passes (or 4 extra function evaluations).

31.3 Exploiting Row (Column) Sparsity

In the discussion of §31.2, we tacitly assume that there are many rows in the Jacobian matrix with ρ nonzeros and that in the new grouping of columns (31.2) each time a structurally orthogonal group is added to a new group $\mathcal{C}_i$, the number of nonzero elements in a row in $\mathcal{C}_i$ is also increased by one. Including more than two structurally orthogonal groups in each of the new groups will reduce the number of new groups and therefore the number of AD passes. Whether the nonzeros in a row of the Jacobian matrix can be recovered by substitution depends on the zero-nonzero structure of the corresponding row of the compressed Jacobian matrix. The compressed Jacobian matrix has p columns, where column j consists of the nonzero elements of the columns in C_j of the Jacobian matrix. Let $p = 8$, $\rho = 5$, and define

$$\mathcal{C}_i = \{C_i, C_{i+1}, C_{i+2}\} \quad i = 1, 2, \ldots, p-2.$$

Further, assume that $\alpha_3 = \alpha_4 = 0$. Then

$$\begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_5 \\ \alpha_6 \\ \alpha_7 \\ \alpha_8 \end{bmatrix} = \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \\ \beta_5 \\ \beta_6 \end{bmatrix},$$

from which the unknowns can be solved for using forward and back substitution. The observation here is that if at least two consecutive elements are known to be zero, then the above system would split and lead to triangular systems.

Consider the matrix (31.4). Each row has at least 2 consecutive zero elements hence the matrix can be computed with 4 passes. In general, if each row of the compressed Jacobian matrix (defined by the consistent partition of columns) has at least d consecutive elements that are known to be zero and $d \leq p - \rho$, then we can define a substitution scheme with $p - d$ passes.

Lemma 2 *If $p > \rho$ and each row of the compressed matrix has at least d consecutive zero entries, $d \leq p - \rho$, then the $p - d$ groups $\{C_1, \ldots, C_{d+1}\}$, $\{C_2, \ldots, C_{d+2}\}, \cdots, \{C_{p-d}, \ldots, C_p\}$ defines a substitution ordering of the nonzero elements.*

All our discussions concerning the new techniques are equally applicable to consistent row partitions and the reverse mode of AD. That is, given a consistent row partition, we define a new grouping of rows that leads to a substitution ordering of the nonzeros, where the nonzeros are obtained from the reverse AD passes.

The above techniques can be extended to the complete direct cover [11] method or bipartition [5]. Given a complete direct cover of the rows and columns of a Jacobian matrix, we can determine the number of consecutive zeros in the compressed Jacobian matrices corresponding to row and column groups and choose the one that gives the largest reduction in the number of AD passes. If the nonzeros in column groups are determined by substitution then the nonzeros in row groups are determined directly and vice versa.

31.4 Computational Results

In this section, we describe the computational experiments for the substitution schemes of §§31.3 and 31.4. Our test problems are drawn from the Harwell-Boeing [7] test matrix collection. For each of the test problems, we compute a row partition, a column partition, and a complete direct cover. We use DSM [3] to compute one directional partitioning and the algorithm in [11] to compute the direct cover. From the sparsity and partition information of the Jacobian matrix, we construct the compressed matrix pattern. For each row of the pattern matrix we calculate the maximum number of consecutive zeros. The minimum of the maximum number of consecutive zeros over all rows (d) is the number of AD passes saved from the substitution ordering. The results on test problems are summarized in Tables 31.1 and 31.2. The following results are presented:

- ρ: Maximum number of nonzeros in any row of A
- mngp: Lower bound on the number of groups in any consistent partition, computed in DSM
- p: Number of groups in the DSM computed consistent partition
- d: Number of consecutive zeros in the compressed matrix obtained from the corresponding consistent partition

TABLE 31.1. Results for SMTAPE collection

	Direct(DSM)						Substitution	
	A			A^T			A	A^T
Name	ρ	mngp	p	ρ	mngp	p	d	d
abb313	6	10	10*	26	26	26	1	0
ash219	2	3	4*	9	9	9	1	0
ash331	2	6	6*	12	12	12	2	0
ash608	2	5	6*	12	12	12	2	0
ash958	2	6	6*	13	13	13	2	0
will199	6	7	7*	9	9	9	1	0
str	34	34	34	34	34	35*	0	1
str200	34	34	34	26	26	31*	0	2
str400	33	33	33	34	36	36*	0	1
str600	33	33	33	34	36	36*	0	1
fs541-1	11	11	13*	541	541	541	1	0
fs541-2	11	11	13*	541	541	541	1	0
lundA	12	12	12	12	12	13*	0	1
lundB	12	12	12	12	12	13*	0	1
Total	200	217	223	1315	1319	1327	11	7

In Table 31.1 we compute consistent partitions for both A and A^T. If we consider the problems where $p > \rho$ (marked with asterisk *) and consider the column computation, then DSM requires 65 matrix-vector products (forward pass of AD), and the substitution technique based on the partition from DSM will reduce the number of matrix-vector products to 54, a reduction of 17%. Even when DSM produces optimal partitioning, e.g., `ash331`, we can reduce the number of AD passes (to 4 from 6).

Our second test suite consists of problems from the CHEMIMP (chemical engineering plant models) collection. We compute a consistent column partition and a complete direct cover for the test problems and report the number of AD passes saved in the corresponding substitution scheme.

TABLE 31.2. Results for CHEMIMP collection

		Direct(DSM)	Substitution	Complete Direct Cover	
Name	ρ	p	d	cgrp(d)	rgrp(d)
impcola	8	8	0	0(0)	6(1)
impcolb	7	11*	2	10(2)	0(0)
impcolc	8	8	0	4(1)	1(0)
impcold	10	11*	1	4(1)	1(0)
impcole	12	21*	1	21(4)	0(0)
Total	45	59	4	39(8)	8(1)

In Table 31.2, cgrp and rgrp are the number of column and row groups respectively in the direct cover. The number of AD passes that can be saved using substitution are enclosed in parenthesis. Taking the sum of the matrix-vector products or AD passes over the problems where $p > \rho$ (marked with asterisk *):

- DSM requires 43 matrix-vector products
- DSM with substitution requires 39 matrix-vector products
- Direct Cover requires 36 AD passes
- Direct Cover with substitution requires 29 AD passes

We get a 19% reduction on the number of AD passes with the Direct Cover technique.

We note that the maximum number of nonzeros in any row or column (ρ) is a lower bound on the number of AD passes needed to determine the Jacobian matrix using either column or row computation. The technique of Newsam and Ramsdell [8, 12] requires only ρ AD passes, and the nonzeros are recovered by solving m small linear systems. With the Vandermonde seed matrices [8], an efficient solver is obtained. However, the numerical reliability of the computed values is a major concern due to ill-conditioning. Other choices, e.g., Chebyshev polynomials for the seed matrices [12], with acceptable condition numbers are possible with an increased computational cost. Our numerical results show that the proposed substitution technique is close to this minimum.

31.5 Concluding Remarks

Our approach to determining the structure of a sparse Jacobian matrix efficiently is based on a given consistent column (or row) partition or a complete direct cover. If the partition is not minimal, the proposed substitution technique saves at least one AD pass. On several test problems, however, we save more than one AD passes.

The number of consecutive zeros in the compressed Jacobian matrix determines how many AD passes we can save. An important research question is whether the number of consecutive zeros in the compressed matrix can be increased by a reordering of the groups or by an interchange of columns (of the Jacobian matrix) between groups.

Acknowledgments: This research was supported in part by the Research Council of Norway and Natural Sciences and Engineering Research Council of Canada under RGPIN.

[1] Brett M. Averick, Jorge J. Moré, Christian H. Bischof, Alan Carle, and Andreas Griewank. Computing large sparse Jacobian matrices using automatic differentiation. *SIAM J. Sci. Comput.*, 15(2):285–294, 1994.

[2] Thomas F. Coleman and Jin-Yi Cai. The cyclic coloring problem and estimation of sparse Hessian matrices. *SIAM J. Alg. Disc. Meth.*, 7(2):221–235, 1986.

[3] Thomas F. Coleman, Burton S. Garbow, and Jorge J. Moré. Software for estimating sparse Jacobian matrices. *ACM Trans. Math. Software*, 10(3):329–345, 1984.

[4] Thomas F. Coleman and Arun Verma. Structure and efficient Jacobian calculation. In Martin Berz, Christian Bischof, George Corliss, and Andreas Griewank, editors, *Computational Differentiation: Techniques, Applications, and Tools*, pages 149–159. SIAM, Philadelphia, Penn., 1996.

[5] Thomas F. Coleman and Arun Verma. The efficient computation of sparse Jacobian matrices using automatic differentiation. *SIAM J. Sci. Comput.*, 19(4):1210–1233, 1998.

[6] A. R. Curtis, Michael J. D. Powell, and John K. Reid. On the estimation of sparse Jacobian matrices. *J. Inst. Math. Appl.*, 13:117–119, 1974.

[7] Iain S. Duff, Roger G. Grimes, and John G. Lewis. Users' guide for the Harwell-Boeing sparse matrix collection (Release I). Technical Report tr/pa/92/86, CERFACS, 1992.

[8] Uwe Geitner, Jean Utke, and Andreas Griewank. Automatic computation of sparse Jacobians by applying the method of Newsam and Ramsdell. In Martin Berz, Christian Bischof, George Corliss, and Andreas Griewank, editors, *Computational Differentiation: Techniques Applications, and Tools*, pages 161–172. SIAM, Philadelphia, Penn., 1996.

[9] Donald Goldfarb and Philippe L. Toint. Optimal estimation of Jacobian and Hessian matrices that arise in finite difference calculations. *Mathematics of Computation*, 43(167):69–88, 1984.

[10] A.K.M. Shahadat Hossain. *On the Computation of Sparse Jacobian Matrices and Newton Steps*. PhD thesis, Department of Informatics, University of Bergen, 1998. Technical Report 146.

[11] A.K.M. Shahadat Hossain and Trond Steihaug. Computing a sparse Jacobian matrix by rows and columns. *Optimization Methods and Software*, 10:33–48, 1998.

[12] Garry N. Newsam and John D. Ramsdell. Estimation of sparse Jacobian matrices. *SIAM J. Alg. Disc. Meth.*, 4(3):404–417, 1983.

[13] Paul E. Plassmann. Sparse Jacobian estimation and factorization on a multiprocessor. In Thomas F. Coleman and Y. Li, editors, *Large-Scale Optimization*, pages 152–179. SIAM, Philadelphia, Penn., 1990.

[14] Michael J. D. Powell and Philippe L. Toint. On the estimation of sparse Hessian matrices. *SIAM J. Numer. Anal.*, 16:1060–1074, 1979.

[15] Trond Steihaug and A.K.M. Shahadat Hossain. Graph coloring and the estimation of sparse Jacobian matrices with segmented columns. Technical Report 72, Department of Informatics, University of Bergen, 1997.