Deflated Decomposition of Solutions of Nearly Singular Systems

Report #225

Tony F. Chan[1]

March 9, 1982

Revised: 9 October 1983 (for Siam J. Numer. Anal.)

[1]Computer Science Department, Yale University, Box 2158, Yale Station, New Haven, CT 06520. The author's work was supported by the Department of Energy under Contract DE-AC02-81ER10996.

Abstract

When solving the linear system Ax = b where A may be nearly singular and b is not consistent with A, one is often interested in computing a *deflated solution,* i.e., an unique solution to a nearby singular but consistent system $A_s\ x_d = \tilde{b}$. Keller [14, 15] has considered deflated solutions with A_s corresponding to setting a small pivot of the LU-factorization of A to zero. Stewart [22] proposed an iterative algorithm for computing a deflated solution with A_s corresponding to setting the smallest singular value of A to zero. Keller's approach explicitly uses submatrices of the LU-factors whereas Stewart's approach is *implicit* in that it only involves solving systems with A. In this paper, we generalize the concept of a deflated solution to that of a *deflated decomposition* which expresses the solution x in terms of x_d and the null vectors of A_s. We treat such decompositions in a uniform framework that includes the approaches of Keller and Stewart and introduce some new deflated solutions based on the LU-factorization. Moreover, we prove some results that relate the different kinds of deflated solutions. In particular, we prove that the difference between one of the LU-based deflated solutions and the SVD-based deflated solution tends to zero as A tends to exactly singular. In addition, we present non-iterative implicit algorithms for computing the LU-based decompositions. Numerical results verifying the accuracy and stability of the algorithms are presented.

1. Introduction

In many numerical problems (for example, in the numerical treatment of nonlinear eigenvalue problems [5, 13, 16, 17, 19, 20], homotopy continuation methods for solving nonlinear systems [1, 9], nearly decomposable Markov chains [21], compartmental models [8] and constrained optimization problems [10]), one often is faced with the problem of having to solve the linear systems of the form

$$Ax = b \tag{1}$$

where it is possible for A to become nearly singular. For simplicity, we shall assume in this paper that A is square and its nullity is at most one. The complete framework generalizes to higher dimensional null spaces but we shall not discuss that here.

For any singular matrix A_s, with a one dimensional null space spanned by u, if x_d is a solution of the consistent system $A_s\ x_d = \tilde{b}$, then $x_d + \mu\ u$ is another solution for *all* scalars μ. Consequently, for the linear system A x = b with A close to A_s, the exact solution x generally

has a large norm. In many applications (see [3] for example), instead of computing x directly, it is preferable to compute the solution x decomposed in the form

$$x = x_d + \eta\, u, \tag{2}$$

where x_d is *purged* of u, for example by requiring that $u^T x_d = 0$. Of course, one would like to compute x_d accurately. However, if one takes the obvious approach and solves the system $A\,x = b$ directly for x (for example, by Gaussian Elimination) and then orthogonalize x with respect to u to get x_d, the solution x one obtains will be dominated by the vector μu and in finite precision, the accuracy in x_d will deteriorate as a result. In this paper, we propose and analyse algorithms for computing such decompositions in a numerically stable manner.

We call x_d a *deflated solution* of (1) and (2) a *deflated decomposition* of the solution x of (1). This decomposition can be viewed as an accurate representation of the solution x of (1) in terms of two parts: a part in the null space u of A_s and a deflated part purged of u. In Section 2, we shall study the question of existence and uniqueness of deflated solutions and deflated decompositions under a general setting.

Depending on how A_s and $\tilde{b}$ are defined in terms of A and b, there may be many different deflated solutions and decompositions. One possibility is to define A_s to be the nearest singular matrix to A in the Frobenious norm and define x_d to be the minimal length least squares solution to the system $A_s x_d = b$. This corresponds to setting the smallest singular value of A to zero in its singular value decomposition (SVD) to obtain A_s and taking $\tilde{b}$ to be the orthogonal projection of b onto the range space of A_s. To avoid computing the SVD of A explicitly, which is usually much more expensive than solving linear systems [2, 11], Stewart [22] gave an *implicit* algorithm for computing x_d which only requires the ability to solve linear systems with A. In Section 3, we present this algorithm and show how it can be extended to compute the deflated decomposition of x. Since the singular vectors required in his algorithm may be inaccurate, Stewart used a form of iterative improvement to refine the deflated solution. We shall show in Section 6 that if the inverse iteration is arranged so that these approximate singular vectors satisfy a simple relationship, then no iterative improvement is necessary.

In Section 4, we define a class of deflated solutions based on a LU-factorization of A with a small pivot, with A_s obtained by changing some elements of A by amounts of approximately the same size as the smallest pivot in the LU factors. These matrices A_s have the property that their left or right null vectors can be determined accurately (to machine precision) by only one back-

substitution. In Section 6, we present implicit iterative algorithms similar to Stewart's for computing these deflated solutions. We analyse the convergence and stability of these algorithms and show that no iterative improvement is necessary.

In Section 5, we prove some results that relate the various deflated solutions. In particular, we show that the difference between one of the LU-based deflated solutions and the deflated solution based on the singular vectors tends to zero as A tends to exactly singular, whereas for the other LU-based deflated solutions this difference tends to something proportional to $v_s^T b$, where v_s is the left null vector of A_s. These results are verified by numerical experiments in Section 7.

We have made no attempt to survey all related work in this area although we would like to mention the work of Peters and Wilkinson [18]. Throughout this paper, upper case Latin letters denote matrices, lower case Latin letters denote vectors and lower case Greek letters denote scalars. We shall use the notation $||.||$ to denote the Euclidean norm, $(u)_k$ to denote the k-th component of the vector u and P_u with $||u|| = 1$ to denote the orthogonal projector $I\text{-}uu^T$.

2. Existence and Uniqueness

Deflated solutions of (1) are solutions to a nearby singular but consistent system derived from (1). All the deflated solutions that we are going to define are solutions to systems of the form:

$$A_s x_d \equiv S A x_d = R b , \tag{3}$$

$$N x_d = x_d , \tag{4}$$

where S, R and N are matrices that are related to A and where $A_s \equiv SA$ is a singular matrix "close" to A, with a one dimensional null space. We shall denote the normalized left null vector of A_s by v_s and the normalized right null vector by u_s. Further, we shall only use matrices N with a one dimensional null space spanned by a normalized left null v_n and a normalized right null vector u_n.

The following lemma gives necessary and sufficient conditions for the existence and uniqueness of the solution to the above system.

Theorem 1: The system (3) and (4) has a solution if and only if $v_s^T R = 0$. The solution is unique if and only if $v_n^T u_s \neq 0$.

Proof: The conditions for existence is simply that the right hand side Rb be in the

range space of A_s. If $v_s^T R = 0$, then all solutions of (3) are of the form $x_d = x_0 + \nu u_s$, where x_0 is any solution of (3) and ν is an arbitrary constant. Condition (4) gives $(v_n^T x_0 + \nu v_n^T u_s) = 0$ from which we can uniquely determine ν if and only if $v_n^T u_s \neq 0$.

Throughout this paper, we shall use matrices S, R and N of a special form so that the conditions in Thoerem 1 are automatically satisfied.

Definition 2: We shall always use S, R and N of the following forms: $S = I - wv_s^T$, $R = I - yv_s^T$ and $N = I - u_s p^T$ with $||v_s|| = 1$, $v_s^T w = 1$, $v_s^T y = 1$ and $u_s^T p = 1$, and $||w||$, $||y||$ and $||p||$ are bounded independently of the singularity of A. Moreover, if A is singular, then v_s must be chosen to be the null vector of A so that $A_s = A$.

Theorem 3: If S, R and N have the forms in Definition 2, then the system (3) and (4) has a unique solution.

Proof: It can be easily verified that the vector v_s in the definition of S is indeed a left null vector of A_s and $v_s^T R = 0$. Moreover, $v_n = p/||p||$ and so $v_n^T u_s = 1/||p|| \neq 0$. Therefore, the conditions of Theorem 1 are satisfied.

Definition 4: Define u_w with $||u_w|| = 1$ and $\theta \geq 0$ to be the unique vector and scalar satisfying $u_w = \theta w$. Similarly define u_y with $||u_y|| = 1$ and $\kappa \geq 0$ to be the unique vector and scalar satisfying $u_y = \kappa y$.

Note that when A is nonsingular, u_w and u_y are just multiples of $A^{-1}w$ and $A^{-1}y$ and θ and κ are normalization constants. When A is singular, $u_w = u_y = u_s$ and $\theta = \kappa = 0$. When A is nearly singular, u_w and u_y are approximate null vectors of A and θ and κ have small absolute values. In fact, it can easily be verified that u_w is the normalized null vector of A_s, i.e. $u_w = u_s$. The choice of N implies that $u_n = u_s$.

If S and R have these specially simple forms, then the solution x of (1) can be easily expressed in terms of the deflated solution x_d and the vectors u_y and u_w.

Theorem 5: Let x_d be the unique solution satisfying (3) and (4) then the following is a solution of (1) :

$$x = x_d + (v_s^T b/\kappa)u_y - (v_s^T A x_d/\theta)u_w. \tag{5}$$

Proof: To show that x given by (5) satisfy $Ax = b$, note that $Ax = Ax_d + (v_s^T b)y - (v_s^T A x_d)w = SAx_d + (I - R)b = Rb + (I - R)b = b$.

We call (5) a *deflated decomposition* of x. When A is nonsingular, (5) can be interpreted as

a decomposition of the unique solution x of (1) into a part spanned by approximate null vectors of A (which could have a large magnitude) and a deflated part whose magnitude remains bounded. When A is singular, it appears that (5) is not defined because of divisions by zero, but it can still have the following interpretation. First note that $v_s^T A x_d = 0$ because $A = A_s$ when A is singular. If b is consistent with A, i.e. $v_s^T b = 0$, then (5) can be intrepreted as x_d plus an *arbitrary* scalar multiple of u_s and therefore still represents solutions to (1). If $v_s^T b \neq 0$, then there is no solution to (1) but (5) can still be interpreted as exhibiting a unique solution of a nearby singular but consistent system (x_d), a null vector of A ($u_s = u_w = u_y$) and the amount that b is inconsistent with A ($v_s^T b$). Thus, the concept of deflated solution and deflated decomposition is meaningful for both singular and nonsingular systems.

With the above interpretation, the deflated decomposition is *unique* in the following sense.

Theorem 6: If the vector $x = z + \rho u_y - (v_s^T A z/\theta) u_w$ with $Nz = z$ satisfies the equation $Ax = b$, then $z = x_d$ and $\rho = v_s^T b/\kappa$.

Proof: Form Ax and it can be verified that it is equal to $A_s z + \rho\kappa y$. But $Ax = b = Rb + (v_s^T b)\, y$. Thus $A_s z = Rb + (v_s^T b - \rho\kappa)y$. If A is nonsingular, then multiplying on the left by v_s^T shows that the coefficient of the last term must be zero and therefore $\rho = (v_s^T b)/\kappa$. If A is singular, then $v_s^T b = 0$ and $\kappa = 0$ implying that ρ may have an arbitrary value. In any case, the last term is zero and it follows that $A_s z = Rb$ and since $Nz = z$, from the uniqueness part of Theorem 3, $z = x_d$.

For the deflated solution to be meaningful, the matrices S and R should be chosen so that A_s is "close" to A and Rb is close to b in some norm. In particular, if A is singular then S is chosen so that $A_s = A$. For computational purpose, the matrix N must also be chosen so that x_d remains bounded independent of the singularity of A. We see from the form of S that $A_s - A = w\,(v_s^T A)$ and therefore it is natural to use approximate left null vectors of A for v_s. In the following two sections, we shall see how v_s can be defined in terms of the SVD of A and the LU-factorization of A. Even with v_s fixed, there are some leeways in choosing the vectors w, y and p and this gives rise to a variety of deflated solutions and deflated decompositions.

3. Deflation Using Singular Vectors

Let u_{sv} and v_{sv} be the normalized right and left singular vectors respectively corresponding to the smallest singular value σ of A. Then we have

$$A\, u_{sv} = \sigma\, v_{sv} \tag{6}$$

and

$$A^T v_{sv} = \sigma\, u_{sv}. \tag{7}$$

Stewart's definition of a deflated solution [22] and our extension to the corresponding deflated decomposition correspond to using $v_s = w = y = v_{sv}$ and $p = u_{sv}$. In other words, one uses $S = R = P_{v_{sv}}$ and $N = P_{u_{sv}}$ in (3) and (4).

Theorem 7: We can write the solution x of (1) as

$$x = x_{sv} + (v_{sv}^T b / \sigma)\, u_{sv} \tag{8}$$

where x_{sv} is the unique solution of the following system:

$$P_{v_{sv}} A\, x_{sv} = P_{v_{sv}} b \tag{9}$$

and

$$P_{u_{sv}} x_{sv} = x_{sv}. \tag{10}$$

Moreover, the matrix

$$\begin{aligned} A_{sv} &\equiv P_{v_{sv}} A \\ &= P_{v_{sv}} A\, P_{u_{sv}} \\ &= A\, P_{u_{sv}} \\ &= A - \sigma\, v_{sv}\, u_{sv}^T \end{aligned} \tag{11}$$

is singular and has v_{sv} and u_{sv} as its left and right null vectors respectively.

Proof: The identities in (11) can be easily verified by using (6) and (7). The existence of a unique deflated solution x_{sv} follows from Theorem 3 and the form of the deflated decomposition follows from Theorem 5.

From Theorem 7, we see that x_{sv} is a deflated solution in the sense that it is a solution to the consistent system (9) with a singular matrix A_{sv} which actually corresponds to setting σ to zero in the SVD of A and is thus the closest singular matrix to A in the Frobenius norm. In fact, by choosing $N = P_{u_{sv}}$, x_{sv} is the minimum length *least squares* solution of the system $A_s x = b$.

4. Deflation Using LU-factorizations

We assume that we can compute a LU-factorization of A (of size n by n) with a small pivot in the k-th position of the form

$$Q_R \, A \, Q_C = L \, U =$$

$$\begin{matrix} + & L_1 & 0 & 0 & + & + & U_1 & u_1 & R_U & + \\ | & & & & | & | & & & & | \\ | & v_1^T & 1 & 0 & | & | & 0 & \epsilon & u_2^T & | \\ | & & & & | & | & & & & | \\ + & R_L & v_2 & L_2 & + & + & 0 & 0 & U_2 & + \end{matrix} \qquad (12)$$

where Q_R and Q_C are row and column permutation matrices, L is unit lower triangular, U is upper triangular and $u_{k,k} = \epsilon$ is a small pivot. We shall assume without loss of generality that the row and column permutations have been pre-applied to A and work with Q_R A Q_C instead of A. Because the dimension of the null space of A is assumed to be at most one, this implies that the matrices L, U_1 and U_2 are well-conditioned, i.e. $||L^{-1}||$, $||U_1^{-1}||$ and $||U_2^{-1}||$ are $o(\epsilon^{-1})$. In other words, we are assuming that the singularity of A reveals itself solely in the smallness of ϵ. For the deflated solutions to be useful, ϵ must be of about the same order of magnitude as the smallest singular value σ of A. We note, however, that although this is often assumed by many people, it is well-known that it does not always hold with the usual pivoting strategies employed in Gaussian Elimination, as the often quoted example of the matrix $\{a_{i,j} = -1$ if $j>i$, $= 1$ if $i=j$, $= 0$ if $i>j\}$ shows [12, 23]. The reader is referred to [4] for a two pass algorithm that is guaranteed to produce such a small pivot for a general matrix.

Based on the LU-factorization, it is easy to define an approximate left null vector for A.

Definition 8: Define v_{lu}, with $||v_{lu}|| = 1$ and $\alpha \geq 0$ to be the unique vector and scalar satisfying

$$A^T v_{lu} = \alpha \, e_k. \qquad (13)$$

Note that if A is nonsingular, v_{lu} and α can be computed by

$$v_{lu} = \alpha \, A^{-T} e_k, \qquad (14)$$

$$\text{and } \alpha = 1 \, / \, || \, A^{-T} e_k \, ||. \qquad (15)$$

The choice of e_k in (13) ensures that $\alpha = O(\epsilon)$, making v_{lu} an approximate left null vector of A.

Lemma 9: $\alpha = O(\epsilon)$.

Proof: It can be shown by directly computing v_{lu} from (12) that $\alpha = \epsilon / ||L^{-T}(0,..,0,1,-U_2^{-T}u_2)^T||$ where the '1' occurs in the k-th position. Since $||L^{-1}||$ and

$||U_2^{-1}||$ are assumed to be $o(\epsilon^{-1})$, we have $\alpha = O(\epsilon)$.

If A is singular, then v_{lu} is the normalized left null vector of A and $\alpha = 0$.

Similar to the deflated solutions based on the SVD of A, one can use v_{lu} to define A_s by choosing $v_s = w = v_{lu}$ or equivalently choosing $S = P_{v_{lu}}$. Note that this choice of w satisfies $v_s^T w = 1$. However, this is not the only way to define A_s. Instead of the orthogonal projector $P_{v_{lu}}$, one can use an *oblique* projector.

Definition 10: For any vector u with $u_j \neq 0$ and $1 \leq j \leq n$, define $E_u^j = I - u\, e_j^T / u_j$.

The matrix E_u^j is singular and has simple left and right null vectors.

Lemma 11: $E_u^j u = 0$, $(E_u^j)^T e_j = 0$, $(E_u^j)^2 = E_u^j$, $((E_u^j)^T)^2 = (E_u^j)^T$.

Note that the projection operator P_u has almost the same properties, except that $P_u^T u = 0$.

Since v_{lu} has norm one, one can always find an index j such that $(v_{lu})_j^{-1} = O(1)$, independent of ϵ. Now one can define S by choosing $v_s = v_{lu}$ and $w = e_j/(v_{lu})_j$, or equivalently choosing $S = (E_{v_{lu}}^j)^T$. Note that this choice of w also satisfies $v_s^T w = 1$.

We now have two possible A_s's.

Definition 12: $A_e \equiv (E_{v_{lu}}^j)^T A$, $A_p \equiv P_{v_{lu}} A$.

Next we define the two corresponding vectors u_w's.

Definition 13: Define u_p and u_e with $||u_p|| = 1$ and $||u_e|| = 1$ and $\gamma \geq 0$ and $\beta \geq 0$ to be the unique vectors and scalars satisfying

$$A\, u_p = \gamma\, v_{lu}, \tag{16}$$

$$A\, u_e = \beta\, e_j, \tag{17}$$

The following lemma shows that v_{lu}, u_e and u_p are related and that β and γ are $O(\epsilon)$.

Lemma 14: $\gamma = \alpha\, (u_p)_k = O(\epsilon)$, $\beta = \alpha\, (u_e)_k \,/\, (v_{lu})_j = O(\epsilon)$.

Proof: Follows easily by left multiplying (16) and (17) by v_{lu}^T and the fact that $\alpha = O(\epsilon)$ and $(v_{lu})_j^{-1} = O(1)$.

It is straigt-forward to prove the following identities for A_e and A_p.

Lemma 15:

$$A_e \equiv (E_{v_{lu}}^j)^T A$$

$$= (E^j_{v_{lu}})^T A E^k_{u_e}$$
$$= A E^k_{u_e}$$
$$= (E^j_{v_{lu}})^T A P_{u_e}$$
$$= A - (\alpha/(v_{lu})_j)\, e_j\, e_k^T \qquad (18)$$

is singular and has v_{lu} and u_e as its left and right null vectors.

$$A_p \equiv P_{v_{lu}} A$$
$$= P_{v_{lu}} A E^k_{u_p}$$
$$= A E^k_{u_p}$$
$$= P_{v_{lu}} A P_{u_p}$$
$$= A - \alpha\, v_{lu}\, e_k^T \qquad (19)$$

is singular and has v_{lu} and u_p as its left and right null vectors.

Thus we see from (18) that A_e is obtained from A by perturbing it at the (j,k)-th position by a quantity that is $O(\epsilon)$. We also see from (19) that A_p is obtained from A by perturbing the k-th column of A by quantities that are $O(\epsilon)$.

Note that the operators $E^k_{u_e}$ and $E^k_{u_p}$ are well defined because of the following Lemma.

Lemma 16: $(u_e)_k \neq 0$, $(u_e)_k \neq 0$.

Proof: When A is nonsingular, this follows from Lemma 14. When A is singular, then $u_e = u_p = u_s$ and direct computation shows that the unnormalized u_s is equal to $(-U_1^{-1}u_1, 1, 0,..,0)$, where the '1' is in the k-th position. Therefore, $(u_s)_k \neq 0$. In fact, since $||U_1^{-1}|| = o(\epsilon^{-1})$, $(u_s)_k^{-1} = O(1)$.

Next, we have to choose R so that the existence condition $v_s^T R = 0$ of Theorem 1 is satisfied. It can easily be seen that the choice of $R = P_{v_{lu}}$ or $(E^j_{v_{lu}})^T$ will do. This corresponds to choosing $y = v_{lu}$ or $e_j/(v_{lu})_j$ and the corresponding u_y's have already been defined in (17) and (16).

Lastly, we have to choose N so that the deflated solution is unique. It turns out that either P_{u_s} or $E^k_{u_s}$ will do. The choice of $N = P_{u_s}$ corresponds to $p = u_s$ and the choice of $N = E^k_{u_s}$ corresponds to $p = e_k/(u_s)_k$. For either choice, $u_s^T p = 1$ and uniqueness follows from Theorem 3.

The choice of $N = P_{u_s}$ corresponds to defining x_d to be the unique minimum length solution to $A_s x = Rb$. The choice of $N = E^k_{u_s}$ makes x the unique solution to $A_s x = Rb$ with $x_k = 0$. Choosing $R = P_{v_{lu}}$ and $N = P_{u_s}$ makes x_d the unique minimum length *least squares* solution to the system $A_s x = b$.

To summarize, we have two possibilities for S (namely $(E^j_{v_{lu}})^T$ and $P_{v_{lu}}$), two for R (namely $(E^j_{v_{lu}})^T$ and $P_{v_{lu}}$), and two for N (namely $E^k_{u_s}$ and P_{u_s}), giving rise to eight deflated solutions and their corresponding deflated decompositions. We shall adopt the following notation.

Definition 17: We shall denote the LU-deflated solutions to (3) and (4) by x_{SRN}, where S, R and N can be either e or p depending on whether the E or the P operator is used. For example, x_{epe} corresponds to using $S = (E^j_{v_{lu}})^T$, $R = P_{v_{lu}}$ and $N = E^k_{u_e}$.

It follows from Theorem 3 that x_{SRN} is well defined.

Theorem 18: The corresponding systems (3) and (4) defining the deflated solutions x_{SRN} have unique solutions.

The corresponding deflated decompositions are given in Table 4-1 and follow directly from Theorem 5.

Some of the deflated solutions based on the LU-factorization have been defined by Keller [14, 15], although he did not use the implicit formulation that we use here. Specifically, he considered the case where the small pivot occurs at the (n,n)-th position of U (i.e. $k = n$). In that case, the unnormalized v_{lu} can be shown to be $(L_1^{-T} v_1 , -1)^T$. If we choose j to be n also, then the unnormalized u_e is $(U_1^{-1} u_1 , -1)^T$. Using these null vectors, he explicitly derived the x_{eee}-deflated decomposition in [15]. In [14], he considered two deflated solutions, which in our framework, correspond to x_{epe} and x_{epp}.

5. Relations Among Deflated Solutions

Each of the deflated solution we have defined so far is the unique solutions to a singular but consistent system derived from (1). Since this singular system is supposed to be close to (1), it is not surprising that some of the deflated solutions are close approximations to one another. In this section, we prove some results that relate the LU-based deflated solutions among themselves and with the deflated solution based on singular vectors (i.e. x_{sv}).

First we shall show that there are certain simple relationships among the LU-based deflated

Table 4-1: Deflated Decompositions of the Solution of $Ax = b$

S	N	R: E	R: P
E	E	x_{eee} + $(v_{lu}^T b / (v_{lu})_j \beta)\ u_e$	x_{epe} + $(v_{lu}^T b / \gamma)\ u_p$
	P	x_{eep} + $[(v_{lu}^T b - \alpha(x_{eep})_k) / (v_{lu})_j \beta]\ u_e$	x_{epp} - $(\alpha(x_{epp})_k / (v_{lu})_j \beta)\ u_e$ + $(v_{lu}^T b / \gamma)\ u_p$
P	E	x_{pee} + $(v_{lu}^T b / (v_{lu})_j \beta)\ u_e$	x_{ppe} + $(v_{lu}^T b / \gamma)\ u_p$
	P	x_{pep} - $(\alpha(x_{pep})_k / \gamma)\ u_p$ + $(v_{lu}^T b / (v_{lu})_j \beta)\ u_e$	x_{ppp} + $[(v_{lu}^T b - \alpha(x_{ppp})_k) / \gamma]\ u_p$

x_{SRN} is the unique solution of $S\,A\,x = R\,b$, $N\,x = x$.

$A^T v_{lu} = \alpha e_k$, $\quad Au_e = \beta e_j$, $\quad Au_p = \gamma v_{lu}$

solutions.

Theorem 19:

(a) $x_{eee} = x_{pee}$ (f) $x_{epe} = x_{ppe}$

(b) $P_{u_e} x_{eee} = x_{eep}$ (g) $P_{u_e} x_{epe} = x_{epp}$

(c) $E^k_{u_e} x_{eep} = x_{eee}$ (h) $E^k_{u_e} x_{epp} = x_{epe}$

(d) $P_{u_p} x_{pee} = x_{pep}$ (i) $P_{u_p} x_{ppe} = x_{ppp}$

(e) $E^k_{u_p} x_{pep} = x_{pee}$ (j) $E^k_{u_p} x_{ppp} = x_{ppe}$

Proof: These equalities can most easily be proven by showing, using (18) and (19), that the difference between the left and right hand sides, say d, satisfies the system (3) and (4) with *homogeneous* right hand sides. Uniqueness then implies d = 0. For example, let $d = E^k_{u_e} x_{eep} - x_{eee}$. Then it can be shown that d satisfies $A_e d = 0$ and $E^k_{u_e} d = d$. The uniqueness result of Theorem 1 implies that d = 0 which proves (c).

We thus see that there are two families of deflated solutions, namely (a) - (e) and (f) - (j), depending on the choice of the R-operator.

Next we will prove that, if ϵ and σ are both small, then v_{lu} is a good approximation to v_{sv}, u_e is a good approximation to u_{sv} and that u_p is an even better approximation to u_{sv}. We shall use the notation $O(\delta_1,\delta_2)$ to mean some quantity (scalar, vector or matrix) the norm of which is bounded by max $\{ O(\delta_1), O(\delta_2) \}$.

Lemma 20: (a) $||v_{lu} - v_{sv}|| = O(\epsilon, \sigma^2)$,

(b) $||u_e - u_{sv}|| = O(\epsilon, \sigma^2)$,

(c) $||u_p - u_{sv}|| = O(\epsilon^2, \sigma^2)$.

Proof: We shall prove (c) first. From the definition of u_p, we have

$$Au_p - \gamma v_{lu} = 0.$$

Now if we multiply both sides of the above equation by A^T, we get

$$(A^T A - \gamma A^T v_{lu} u_p^T)\, u_p \equiv M_1\, u_p = 0.$$

Similarly, for the singular vectors u_{sv} and v_{sv}, we have

$$(A^T A - \sigma A^T v_{sv} u_{sv}^T]\, u_{sv} \equiv M_2\, u_{sv} = 0.$$

Thus both u_p and u_{sv} are null vectors to matrices that are perturbations of the symmetric $A^T A$. We also have

$$\begin{aligned} M_1 &= A^T A - \gamma\alpha e_k u_p^T \\ &= M_2 + \sigma^2 u_{sv} u_{sv}^T - \gamma\alpha e_k u_p^T \\ &= M_2 + O(\sigma^2, \epsilon^2). \text{ (from Lemma 9)} \end{aligned}$$

Therefore, M_1 is a $O(\sigma^2, \epsilon^2)$ perturbation of M_2. From a standard perturbation result for simple eigenvalues ([23], p.67), (c) follows immediately.

The proof for (b) is almost identical, except that now

$$\begin{aligned} M_1 &= M_2 + \sigma^2 u_{sv} u_{sv}^T - \beta A^T e_j u_p^T \\ &= M_2 + O(\sigma^2, \epsilon). \end{aligned}$$

because although β is $O(\epsilon)$, $A^T e_j$ is $O(1)$ in general. The proof for (a) is similar.

Next we use Lemma 20 to prove the main result of this section.

Theorem 21: If A is nonsingular, then

(a) $||x_{eep} - x_{sv}|| \leq O(\sigma^2,\epsilon) (|v_{sv}^T b/\sigma| + c_1||x_{eep}||)$

(b) $||x_{ppp} - x_{sv}|| \leq O(\sigma^2,\epsilon^2) (|v_{sv}^T b/\sigma| + c_2||x_{ppp}||)$

where c_1 and c_2 are positive constants independent of ϵ and σ. If A is singular, then $x_{ppp} = x_{sv}$ and $x_{eep} - x_{sv} = (v_{sv}^T b)\, A^\dagger q$, where $A^\dagger$ is the pseudo-inverse of A and $q = v_{sv} - e_j/(v_{sv})_j$.

Proof: We shall start with the nonsingular case and prove (a) first. The strategy is to take the x_{sv}-decomposition (8) and orthogonalize the two parts with respect to u_e. Since the x_{eep}-decomposition is a unique decomposition of x, by Theorem 6 the parts that are orthogonal to u_e must be equal to x_{eep}. Thus, we have

$$x_{sv} = [x_{sv} - (u_e^T x_{sv})u_e] + (u_e^T x_{sv})u_e ,$$

$$u_{sv} = [u_{sv} - (u_e^T u_{sv})u_e] + (u_e^T u_{sv})u_e .$$

It follows that, with $\nu = v_{sv}^T b/\sigma$,

$$x_{eep} = [x_{sv} - (u_e^T x_{sv})u_e] + \nu[u_{sv} - (u_e^T u_{sv})u_e].$$

From Lemma 20, part (b), we can easily deduce that

$$u_e^T x_{sv} = (u_{sv} + O(\epsilon,\sigma^2))^T x_{sv} = O(\epsilon,\sigma^2) x_{sv} , \tag{20}$$

and

$$u_{sv} - (u_e^T u_{sv}) u_e = u_e + O(\epsilon,\sigma^2) - (1 + O(\epsilon,\sigma^2)) u_e = O(\epsilon,\sigma^2) , \tag{21}$$

from which (a) follows immediately. The proof for (b) is analogous, except that because of Lemma 20, part (c), we have $O(\epsilon^2,\sigma^2)$ in (20) and (21). When A is singular, $v_{lu} = v_{sv}$ and therefore x_{sv} and x_{ppp} are defined by the same S, R and N. The uniqueness result of Theorem 3 implies that they must be identical. Let $d = x_{eep} - x_{sv}$. Then it can be shown that d satisfies $A\,d = (v_{sv}^T b)\,q$ and $P_{u_{sv}}\,d = d$. It follows that d is the minimum length least squares solution and therefore $d = (v_{sv}^T b)\,A^\dagger\,q$.

The above theorem implies, if $\epsilon = O(\sigma)$, that x_{ppp} will approach x_{sv} as σ goes to zero but that x_{eep} will generally be different from x_{sv} (unless $v_{sv}^T b = 0$) in the same limit.

6. Algorithms

In this section, we propose *implicit* algorithms for computing the deflated solutions and deflated decompositions defined in Section 4 and analyse their convergence and stability properties. It should be apparent that the primary task is to compute the deflated solutions because the deflated decompositions can then be computed without too much difficulty. Since the implicit algorithms use as a basic tool the ability to solve systems with A, we shall assume in this section that A is *nonsingular relative to the precision of the computer*. In practice this is almost always true. We shall limit our discussions to direct methods based on Gaussian Elimination.

First of all, we have to compute the approximate left and right null vectors of A_s. For v_{lu}, u_e and u_p, they can each be computed by one backsubstitution using the LU factors of A by formulas similar to (14) and (15). For v_{sv} and u_{sv}, we can use an inverse iteration similar to one proposed by Stewart [22]:

Inverse Iteration for v_{sv}, u_{sv} and σ

Starting with an initial guess v_0, iterate until convergence:

1. $A\tilde{u}_{i+1} = v_i$
2. $u_{i+1} = \tilde{u}_{i+1}/||\tilde{u}_{i+1}||$
3. $A^T\tilde{v}_{i+1} = u_{i+1}$
4. $v_{i+1} = \tilde{v}_{i+1}/||\tilde{v}_{i+1}||$.

Denote the converged v_i by v_{sv}, compute u_{sv} and σ by:

$$A\tilde{u} = v_{sv},$$

$$u_{sv} = \tilde{u}/||\tilde{u}||,$$

$$\sigma = 1 \;/\; ||\tilde{u}||. \tag{22}$$

After the null vectors of A_s have been computed, one then has to find algorithms for solving the system (3) and (4) for the deflated solutions. We propose the following algorithm for computing the unique solution x_{SRN} of the system (3) and (4) which is based on a similar algorithm first proposed by Stewart [22].

Algorithm IIA (Iterative Improvement Algorithm)

Start with any x such that $Nx = x$

Loop until convergence:

(1) $r = Rb - SAx$

(2) Solve $Ad = r$

(3) $x \Leftarrow x + Nd$

The following theorem states the conditions under which Algorithm IIA will converge to the desired solution x_{SRN}.

Theorem 22: Assume that the assumptions of Theorem 1 are satisfied so that the system defining x_{SRN} has a unique solution. Further assume that N has a one dimensional null space u_n and that $N^2 = N$. Define $K \equiv I - A^{-1}A_s$, $M \equiv NKN$ and $\mu \equiv v_s^T A u_n$. Then the following statements are true for Algorithm IIA:

(a) All iterates x satisfy $Nx = x$.

(b) If x converges, then it will converge to x_{SRN} if $\mu \neq 0$.

(c) The iterate x converges if $||M|| < 1$.

(d) If $M = 0$, then we can obtain x directly from $x = NA^{-1}Rb$.

Proof: Since we always start with an x such that $Nx = x$, Step (3) of Algorithm IIA guarantees that all iterates satisfy the same constraint since $N^2 = N$. The iteration can be written as the following linear stationary iteration:

$$x \Leftarrow (I - NA^{-1}A_s)x + NA^{-1}Rb. \tag{23}$$

If the iteration converges to x, then we have $x = (I - NA^{-1}A_s)x + NA^{-1}Rb$ from which follows that $NA^{-1}r(x) = 0$, where $r(x) = Rb - A_s x$. This implies that $r(x) = \nu\, Au_n$, where ν is an arbitrary scalar. Left multiplying by v_s^T, we see that if $\mu \neq 0$, then $\nu = 0$ and therefore $r(x) = 0$, which together with $Nx = x$ proves that x satisfies (3) and (4). Uniqueness then implies $x = x_{SRN}$. For analysing the convergence, note that since $Nx = x$ and $N^2 = N$, we can rewrite the iteration as

$$x \Leftarrow N(I - A^{-1}A_s)Nx + NA^{-1}Rb = Mx + NA^{-1}Rb, \tag{24}$$

and thus the iteration will converge if $||M|| < 1$. If $M = 0$, then we have convergence after one iteration independent of the starting vector and the converged solution is $NA^{-1}Rb$ which must be equal to x_{SRN} because of uniqueness.

We can now apply the above theorem to the application of Algorithm IIA for computing x_{sv} and the x_{SRN}'s. The matrix K only depends on S, whereas M depends on both S and N. Since we always choose N to be either $E_{u_s}^k$ or P_{u_s}, μ only depends on the choice of S. In Table 6-1, we give the expressions for K, M and μ for the different possible choices of S and N.

From Table 6-1, we see that all the M's are exactly equal to zero and all the μ's are nonzero. Thus, for all the deflated solutions that we have considered so far, Algorithm IIA will converge in *one* step, and consequently we can use the non-iterative version as outlined in Theorem 22, part (d).

The above conclusion assumes that we have the exact null vectors v_s and u_s of A_s available. Although the LU-based null vectors v_{lu}, u_e and u_p can most likely be computed with small relative errors, the accuracy of v_{sv} and u_{sv} depends on the convergence of the inverse iteration. If the smallest singular value of A is not well isolated, then the inverse iteration may have

Table 6-1: Table of K, M and μ

		M		
S	K	$N = E^k_{u_s}$	$N = P_{u_s}$	μ
$P_{v_{sv}}$	$u_{sv}u^T_{sv}$	N.A.	0	σ
$(E^k_{v_{lu}})^T$	$u_e e^T_k/(u_e)_k$	0	0	$(u_e)_k$
$P_{v_{lu}}$	$u_p e^T_k/(u_p)_k$	0	0	$(u_p)_k$

convergence difficulty. This may occur if one uses the deflation algorithm when A is not nearly singular. If v_{sv} and u_{sv} are completely unrelated and do not satisfy (6) and (7), then $M \neq 0$ and Algorithm IIA will in general take more than one iteration to converge. However, in the implementation of the inverse iteration in (22), v_{sv} and u_{sv} do satisfy (6). It turns out that this is enough for M to be equal to zero although the deflated decomposition has to be modified because the last term in (5) is no longer equal to zero.

Theorem 23: If the approximate singular vectors u_{sv} and v_{sv} used in Algorithm IIA for x_{sv} satisfy $A\, u_{sv} = \sigma\, v_{sv}$ (but not necessarily $A^T v_{sv} = \sigma\, u_{sv}$), then $M = 0$, and the corresponding deflated decomposition is given by:

$$x = x_d + (v_{sv}^T(b - Ax_d)/\sigma)\, u_{sv}. \tag{25}$$

Proof: The proof is straight-forward and follows from Theorem 5.

Seen in this light, the LU-based x_{ppp} deflated decomposition can also be considered as a member of the SVD-based deflated decomposition where v is chosen specifically to be v_{lu}, or equivalently, using one step of inverse iteration with $v_0 = e_k$.

In choosing N, the only necessary condition is to satisfy the uniqueness condition of Theorem 3. We have chosen N with a null vector u_n equal to u_s. Although this choice is not necessary, we shall argue that it leads to stable algorithms for computing them in finite precision arithmetic. It is well-known that when one performs a back-substitution with an ill-conditioned A, in general the solution will have large errors, and the residual will be large. However, the standard round-off error analysis also shows that if the computed solution is not large, then the residual must be small ([6], p.181), even if A is ill-conditioned. This is exactly what happens here. Step (1) of Algorithm IIA changes the right hand side so that the solution obtained in Step (2) will be small. However, if A is nearly singular, a small *residual* still allows for a possible large *error* in the solution in the direction of the null vector. Therefore it is a good idea to choose N so that its null vector u_n is the same as the null vector u_s of A_s so that Step (3) annihilates this error. This makes the non-iterative algorithm stable to round-off errors. These nice properties will not hold if we use an arbitrary N with $u_n \neq u_s$.

Because of the equivalence relationships in Theorem 19, we can limit our attention to computing only two of the LU-based deflated solutions, corresponding to the two possibilities for R. The other deflated solutions can be obtained from these two by simple transformations. We recommend computing either x_{eee} or x_{eep}, and either x_{ppe} or x_{ppp} because first, the corresponding

deflated decompositions requires computing only two approximate null vectors instead of three for the others, and second, Step (1) of Algorithm IIA can be simplied to S(b-Ax) since S = R.

We now say a few words about the efficiency of the algorithms. We first note that the P-operators take one inner product to apply whereas the E-operators take none. The non-iterative version of Algorithm IIA costs one LU-factorization of A, one back-substitution and the cost of obtaining the null vectors. It also requires storage for the (2 or 3) approximate null vectors and, in the case of (25), storing a copy of A. The iterative version also requires storing a copy of A. Since the factorization usually requires much more time (storage) than the back-substitutions (solution), the work (storage) involved is usually not much more than the normal factor-solve process for (1). If there are more than one righ hand sides, the cost of the extra back-solves for the null vectors can also be amortized over the total computing time. For the LU-based deflations, this extra cost is always two back-solves, regardless of whether A is nearly singular or not. On the other hand, the cost and convergence of the inverse iteration for computing the singular vectors of A are much more sensitive to the singularity of A (two backsolves per iteration). Moreover, if the singular vectors are not accurate, then the computed x_{sv} will not be the minimum length least squares solution to Ax = b (although it will still be a deflated solution in our interpretation), and thus, in view of the results of Theorem 21, it is no more special than the LU-deflated solutions. Furthermore, an extra copy of A must be stored to obtain the deflated decomposition. Therefore, in applications where A may occasionally not be nearly singular or when it is not known a priori whether A is nearly singular or not, we argue that the LU-deflations are to be preferred because they are non-iterative. Such situations arise, for example, in applying continuation methods to solving nonlinear equations with Jacobian matrices that may become singular [3, 13, 16]. If it is known that A is very nearly singular, then the SVD-deflations are probably to be preferred because they do not depend on the assumption that the LU-factorization of A has a small pivot and because the inverse iteration will converge very quickly in that case. Both of these implicit deflation techniques are to be preferred to the explicit deflation techniques if the data structures for storing the LU-factorization of A are complicated, for example, in band solvers and sparse solvers. An extra advantage of the implicit algorithms is its modularity - it is independent of how the factorization and solve is done and requires minimum modifications to the conventional factor-solve procedure. The most effective approach may be in the form of a hybrid algorithm which uses a LU-based algorithm as a default and switches to a SVD-based algorithm when the condition $\epsilon = O(\sigma)$ is not satisfied.

7. Numerical Results

We present some numerical results to verify the accuracy and stability of the various deflation algorithms developed in the earlier sections. We consider two classes of matrices:

A_1 : $(I-2uu^T)$ Diag(σ,n-1,n-2,.....,1) $(I-2vv^T)$ where u and v are chosen randomly and scaled to have norm 1, and σ varies from 1 to 10^{-8}.

A_2 : $T - \lambda_{min}(T)I - \sigma I$ where T = Tridiagonal(1,-2,1) and σ again varies from 1 to 10^{-8}.

Note that the smallest singular value of A_1 and A_2 is equal to σ. For A_1, the smallest singular value has multiplicity 2 when $\sigma = 1$. The dimension n of A is chosen to be 20.

We will only be concerned with computing x_{sv}, x_{eee}, x_{eep}, x_{ppe} and x_{ppp}. Solutions are chosen to have the form $x = z + \rho u_s$ where z is randomly chosen and satisfies $Nz = z$ and u_s is the corresponding null vector of A_s. The right hand side b is then obtained by forming $b = A v + \rho A u_s$ where the last term is formed by using the definitions of u_s in (6), (16) and (17) so as to minimize round-off errors. By Theorem 6, z is the unique deflated solution of $A x = b$. The constant ρ is used to control the value of $v_s^T b$.

For comparison, we will consider the following deflated solution $x_{ge} = P_{u_{sv}} A^{-1} b$, which is equal to x_{sv} in exact arithmetic but computed without using deflation. All LU-factorizations are performed by the routine SGECO of LINPACK [7] which uses the partial pivoting strategy. The computations were performed on a DEC-20 with 27 bits mantissas corresponding to a machine precision of about $.4 \times 10^{-8}$.

The first set of tests is to see how ϵ varies with σ and to verify the accuracy of the computed σ. In the inverse iteration for determining the singular vectors we always take 5 iterations. When A is highly singular, one or two iterations is enough for full accuracy. The computed ϵ, its position k, and the computed σ are given in Table 7-1 for A_1 and A_2. We see that, at least for these two classes of matrices, ϵ is indeed generally $O(\sigma)$. However, the smallest pivot does not always appear at the (n,n)-th position, especially when ϵ is not small. In fact, for the case $\sigma = 10^{-3}$ for A_1, the n-th pivot is not small at all. Note also that, when σ is well-isolated, the computed σ's are rather accurate and have low *absolute* errors. However, when the smallest singular value is not well isolated, the inverse iteration is not successful at all. This is especially true for A_2 because its lowest eigenvalues are rather close to each other.

The next set of tests is to check the accuracy and stability of Algorithm IIA. For each choice of S, the right hand sides b are generated as discussed above so that $v_s^T b = 1$, and the deflated solutions x_{eee}, x_{ppe}, x_{sv} and x_{ge} are computed. The iterative version of Algorithm IIA was implemented but one iteration always proved to be enough for all the deflated solutions, so the results given here are computed with the non-iterative version. The relative errors are displayed in Figs. 7-1 and 7-2. As expected, x_{ge} loses accuracy as A becomes more singular whereas the other deflated solutions remain accurate to within roughly an order of magnitude of the machine round-off level. Moreover, one can draw a distinct correlation between the less accurate LU-deflated solutions with relatively large values of ϵ (e.g. $\sigma = 10^{-5}$ for A_1). These tests indicate that the LU-deflated solutions and the non-iterative version of Algorithm IIA together form a rather robust procedure for computing deflated solutions.

The next set of tests is to verify the results of Lemma 20 and Theorem 21. A large number of matrices of the form A_1 are generated and the three null vectors u_{sv}, u_e and u_p are computed. In Fig. 7-3, $||u_{sv} - u_e||$ and $||u_{sv} - u_p||$ are plotted against σ. It is seen that u_p is indeed generally closer to u_{sv} than u_e is for the same value of σ. The straight lines shown are the best least squares fit to the two sets of data {log(difference), log(σ)} by straight lines. From Lemma 20, the exact slopes of the straight lines should be 1 for $||u_{sv} - u_e||$ and 2 for $||u_{sv} - u_p||$. The slopes of the least squares fit are 1.028 and 1.537 respectively.

Next, right hand sides are generated with values of $v_{sv}^T b$ varying from 10^{-4} to 10^{-1} for the matrices A_1 and all the deflated solutions computed with the same right hand sides. In Figs. 7-4 and 7-5, $||x_{sv} - x_{eep}||$ and $||x_{sv} - x_{ppp}||$ are plotted as functions of σ. It is seen that $||x_{sv} - x_{eep}||$ does tend to a constant value roughly proportional to $v_{sv}^T b$ and $||x_{sv} - x_{ppp}||$ does tend to zero (or round-off level) as σ tends to zero. Moreover, the differences do vary linearly with $v_{sv}^T b$ when σ is small. The results of Lemma 20 and Theorem 21 are thus verified.

8. Conclusion

We have offerred a rather complete analysis of deflated solutions and deflated decompositions of solutions to nearly singular linear systems. We have provided a uniform framework through which it is possible to relate the various approaches used in the literature as well as some new approaches proposed here. We have analysed both theoretical questions of existence and uniqueness and practical questions of stable computational algorithms. The use of

Table 7-1: ϵ as a Function of $\sigma = 10^{-I}$

Matrix A_1

I	ϵ	k	computed σ	
0	0.1504303E+01	1	0.1000000E+01	**
1	0.1130482E+00	1	0.1000000E+00	
2	0.8422814E-01	1	0.1000000E-01	
3	-0.7856242E-01	19	0.9999483E-03	
4	-0.8268356E-03	20	0.9996640E-04	
5	0.9228393E-02	20	0.9993521E-05	
6	-0.1059374E-04	20	0.1034953E-05	
7	-0.7590279E-06	20	0.7292485E-07	
8	-0.2421439E-06	20	0.2352214E-07	

Matrix A_2

I	ϵ	k	computed σ	
0	0.2773407E+00	20	0.2233835E-01	*
1	0.9002221E+00	20	0.7668735E-01	*
2	0.5978710E+00	20	0.1000000E-01	
3	0.6968676E-01	20	0.9999985E-03	
4	0.7055007E-02	20	0.9999866E-04	
5	0.7061549E-03	20	0.9996851E-05	
6	0.7033348E-04	20	0.9955706E-06	
7	0.6996095E-05	20	0.9902853E-07	
8	0.6780028E-06	20	0.9597004E-08	

* Inverse iteration had not converged after 5 iterations.
** Singular vectors had not converged after 5 iterations.

Figure 7-1: Relative Errors of Computed x_{SRN} vs $\sigma = 10^{-I}$ for A_1

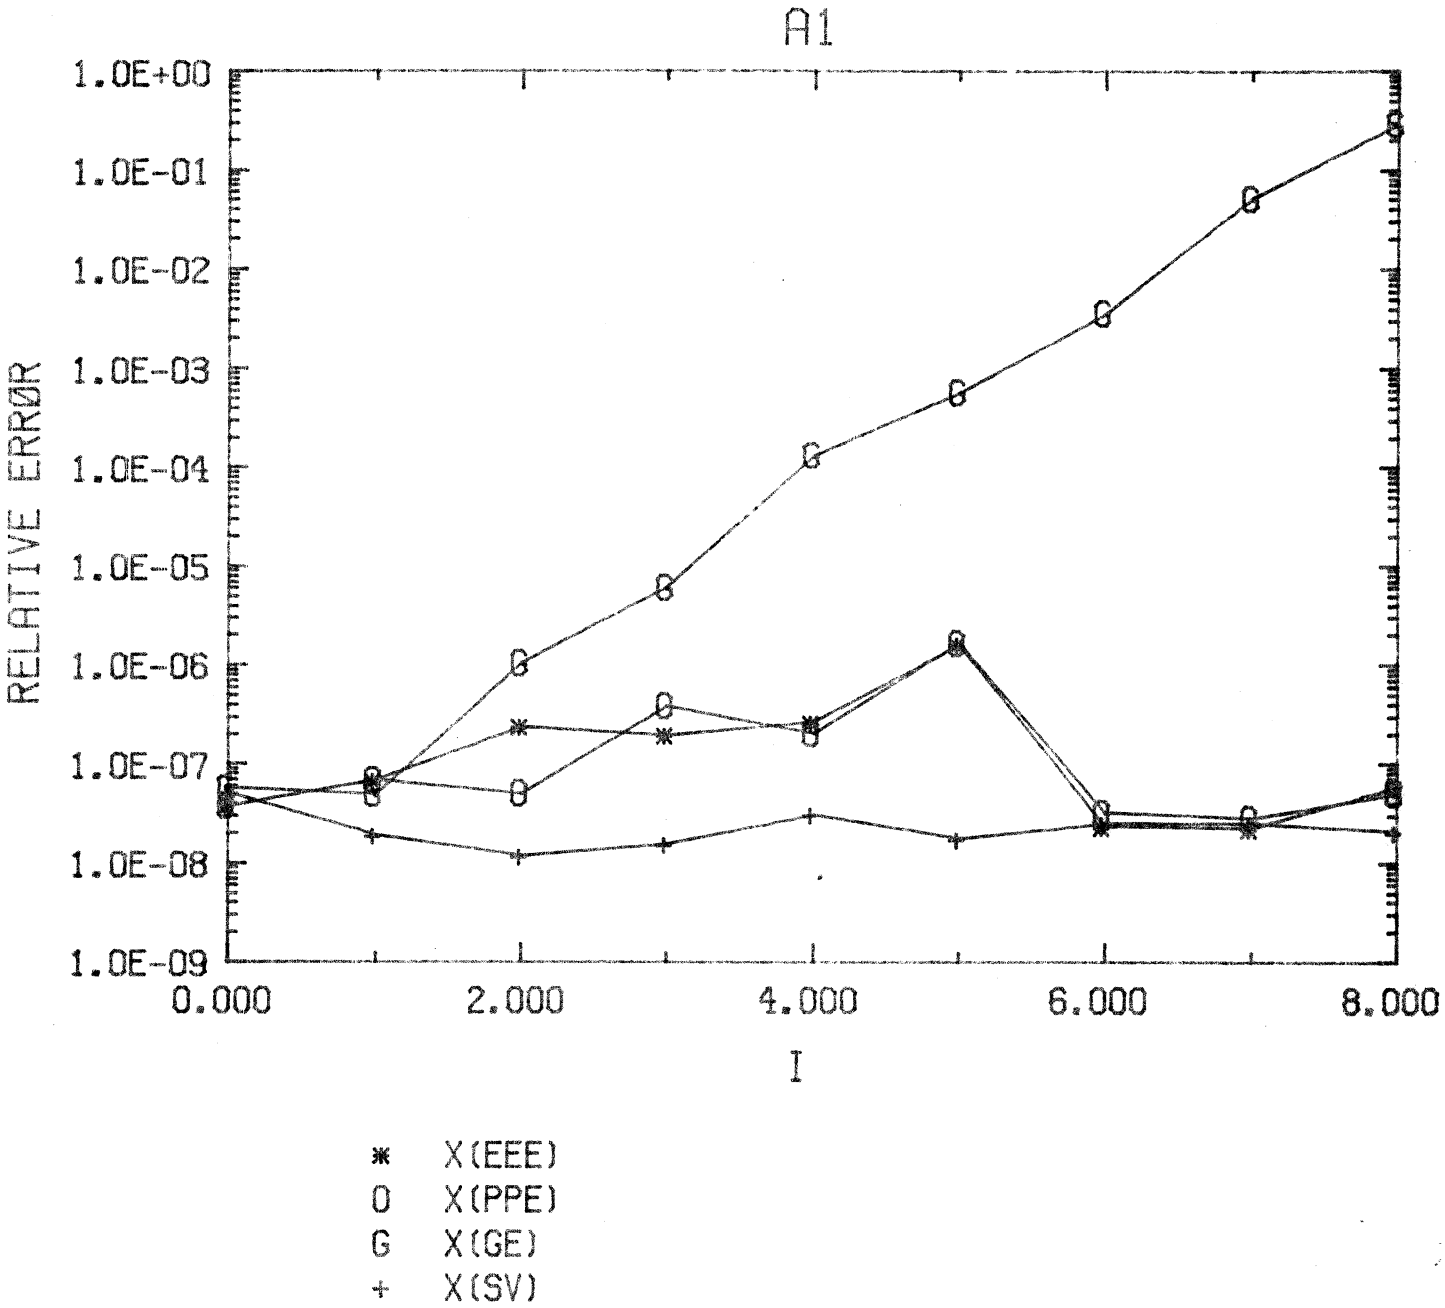

Figure 7-2: Relative Errors of Computed x_{SRN} vs $\sigma = 10^{-I}$ for A_2

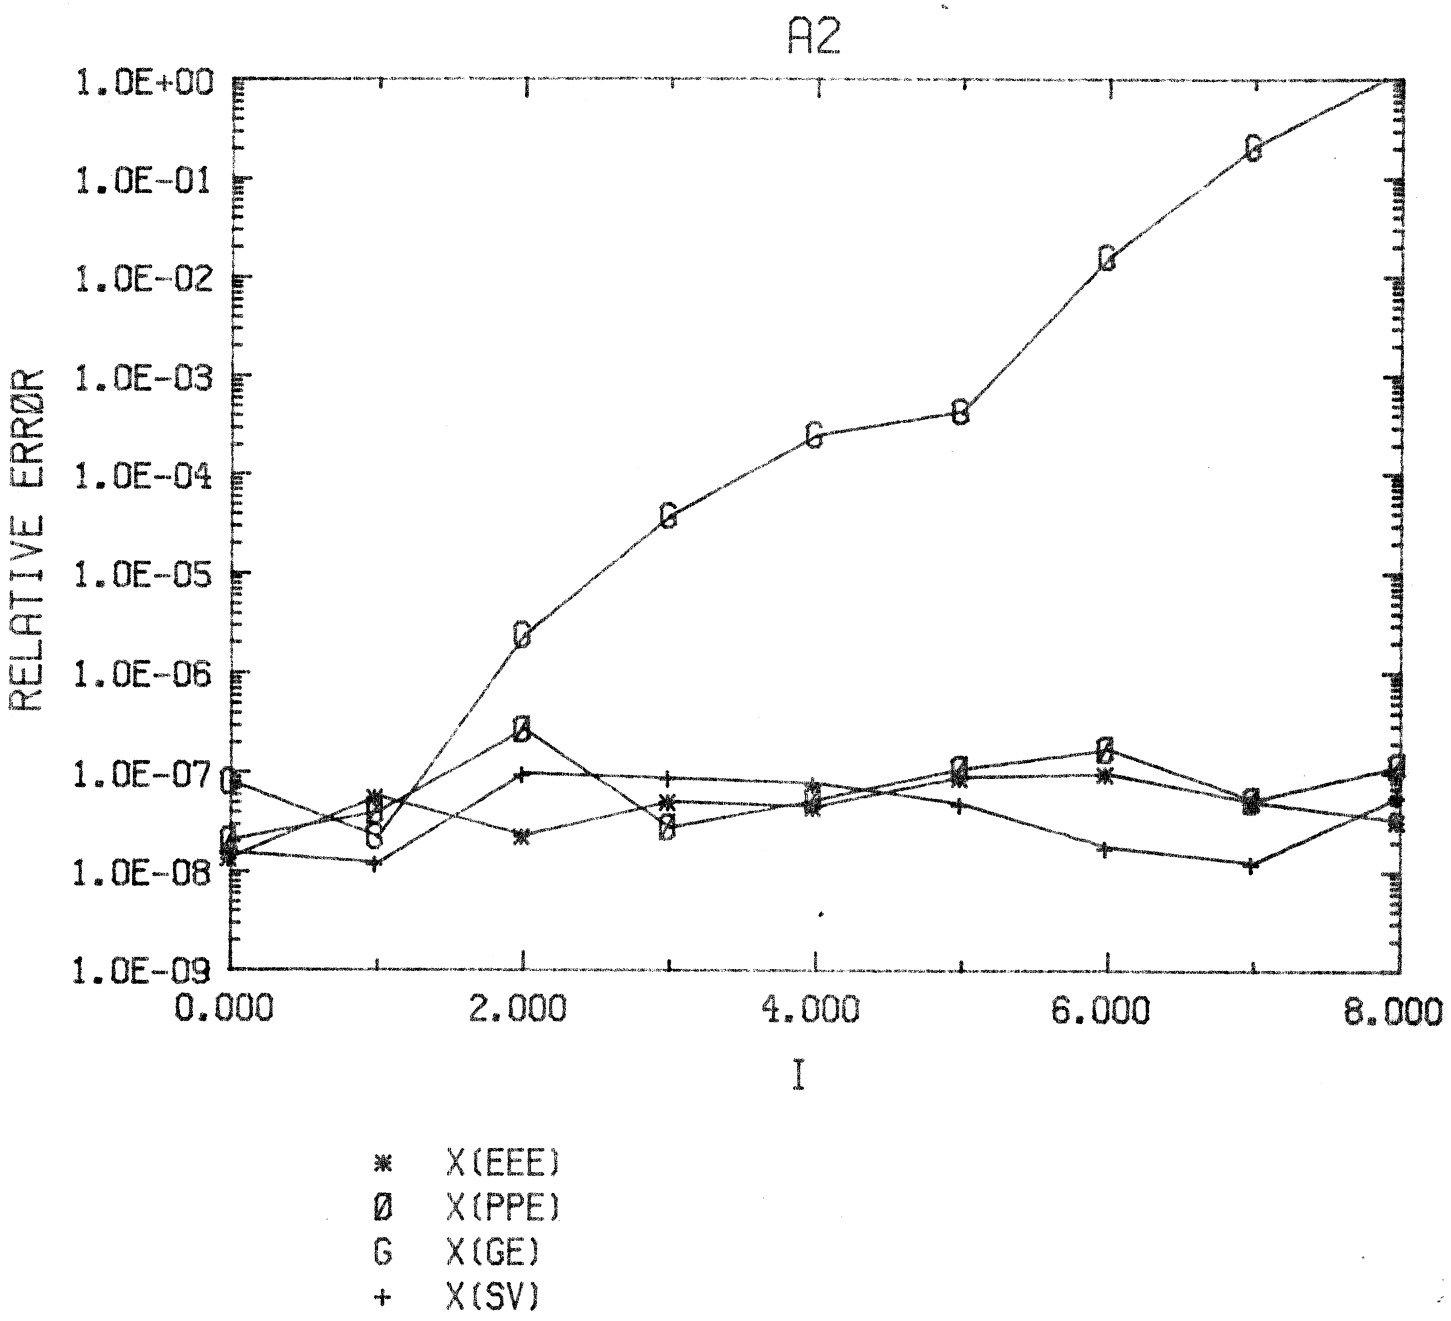

Figure 7-3: $||u_{sv} - u_e||$ and $||u_{sv} - u_p||$ vs. $\sigma = 10^{-I}$

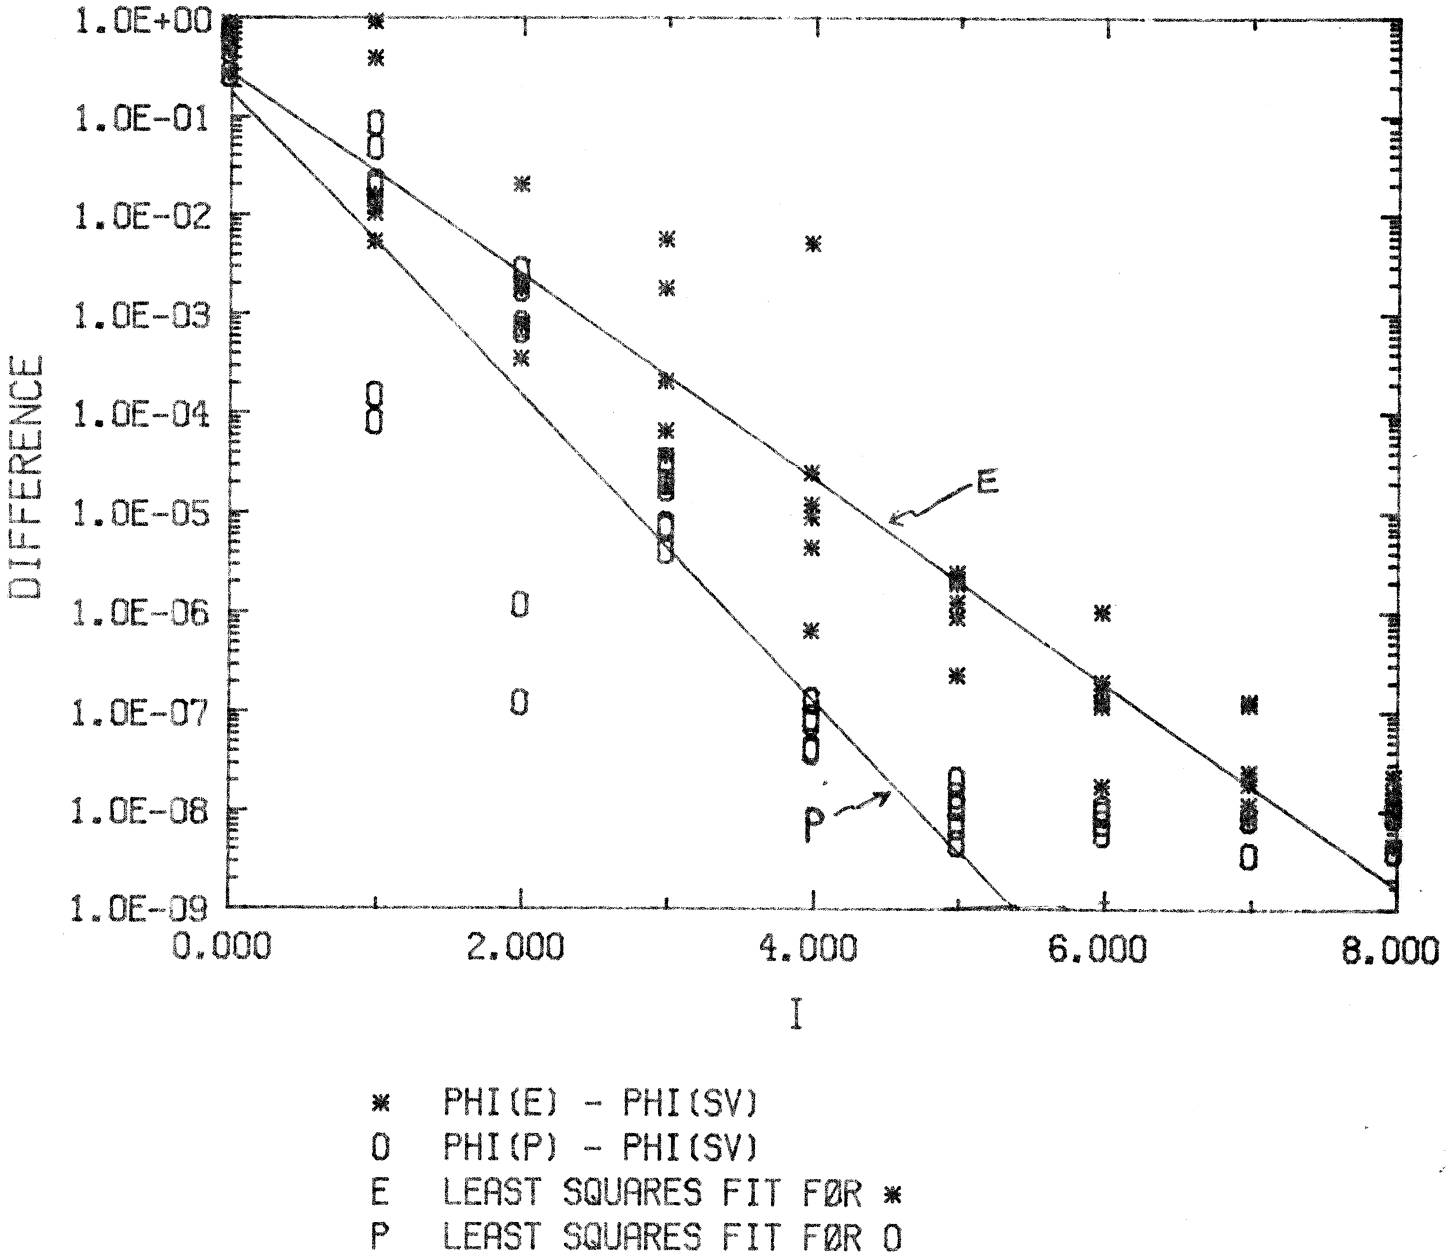

Figure 7-4: $||x_{sv} - x_{eep}||$ vs. $\sigma = 10^{-I}$, $(P = v_{sv}^T b)$

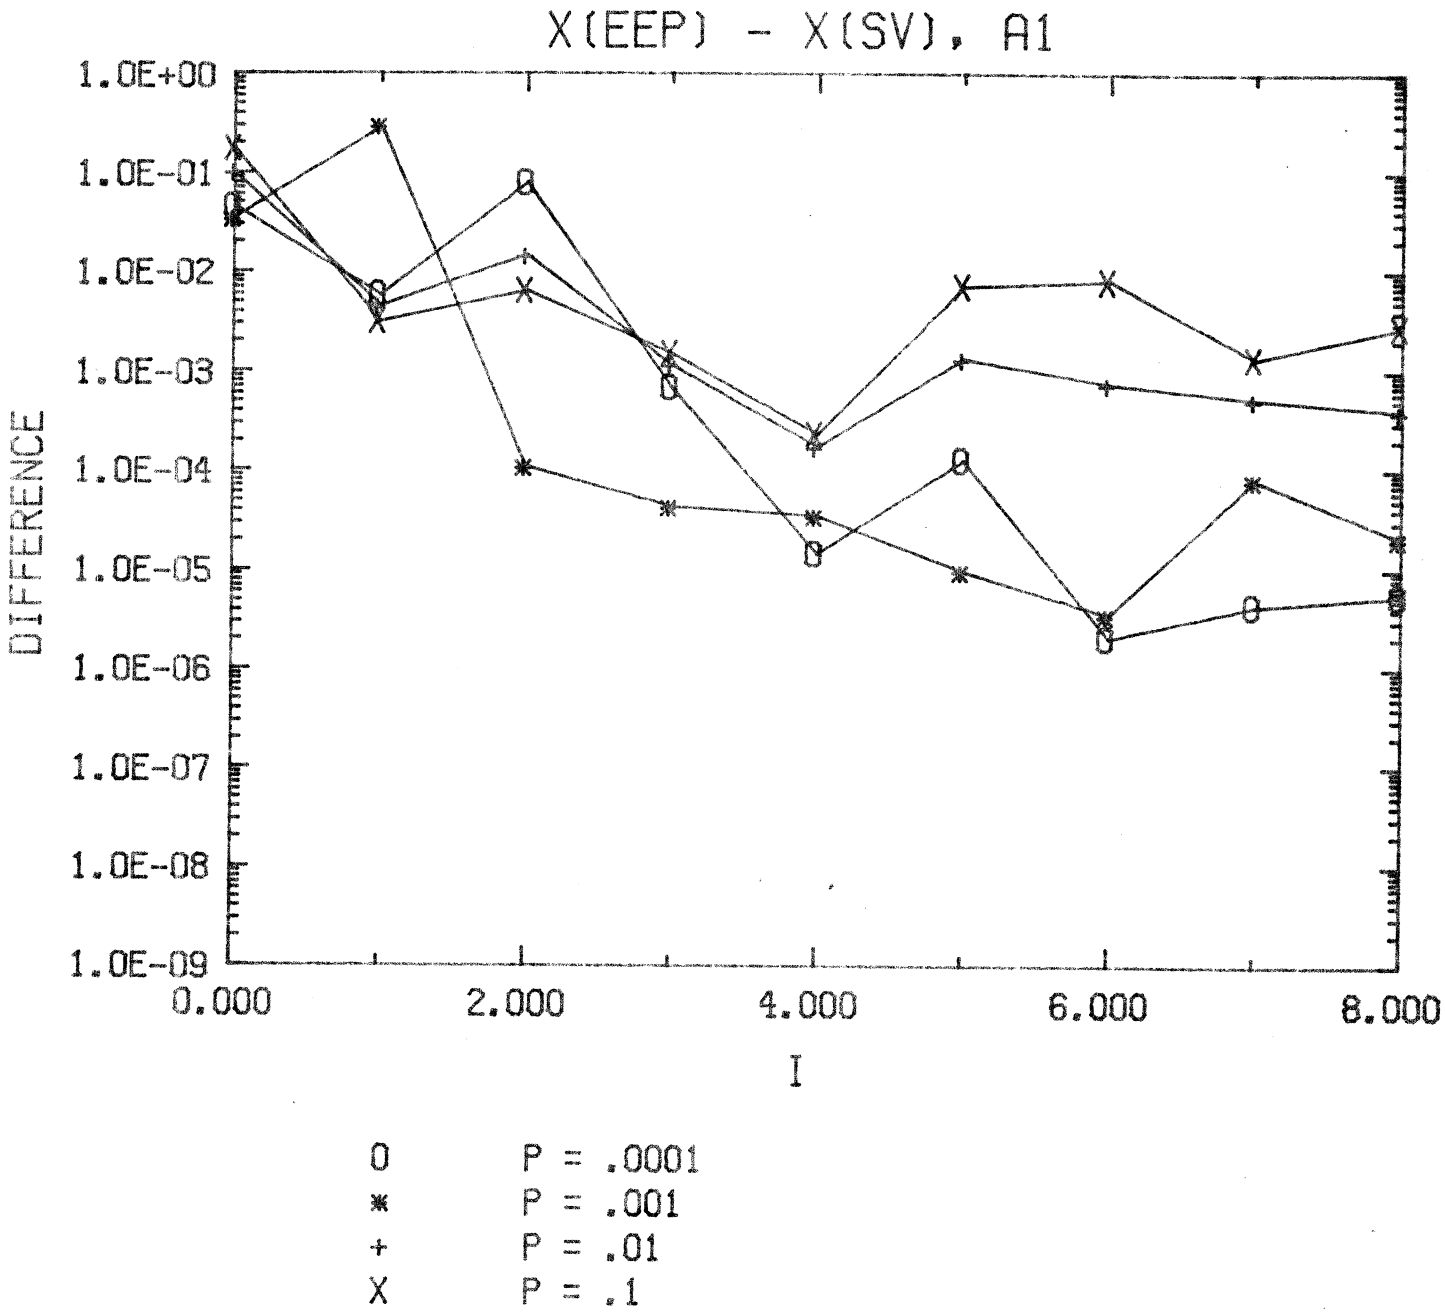

Figure 7-5: $||x_{sv} - x_{ppp}||$ vs. $\sigma = 10^{-I}$, $(P = v_{sv}^T b)$

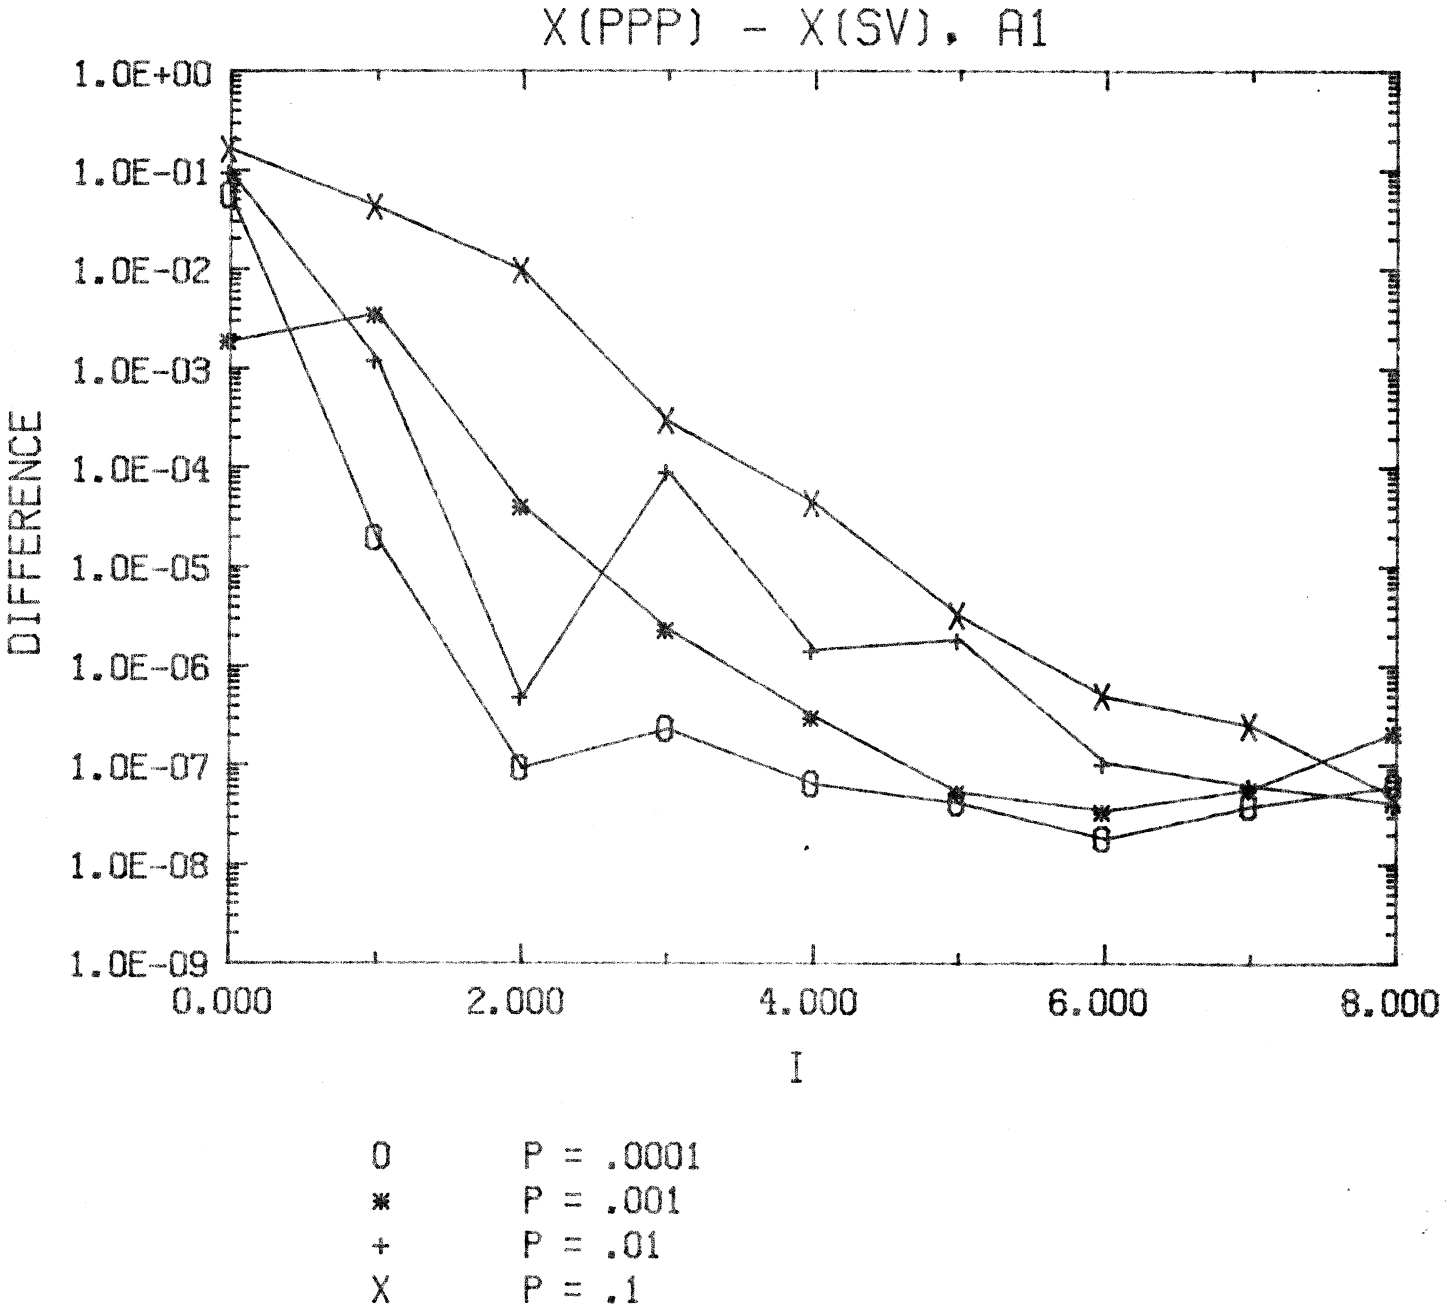

implicit algorithms results in a modular approach which only accesses the matrix as a linear solver. Therefore, the algorithms proposed here should be easily extensible to linear solvers other than dense Gaussian Elimination, for example, sparse direct solvers, conjugate gradient methods and multi-grid methods. Extensions to higher dimensional null spaces should also be straightforward.

Acknowledgement The author would like to thank the editor and one anonymous referee for many helpful suggestions on the presentation.

References

[1] E. Allgower and K. Georg, *Simplicial and Continuation Methods for Approximating Fixed Points and Solutions to Systems of Equations,* SIAM Review, 22 (1980), pp. 28 - 85.

[2] T.F. Chan, *An Improved Algorithm for Computing the Singular Value Decomposition,* ACM TOMS, 8 (1982), pp. 72-83.

[3] T.F. Chan, *Deflation Techniques and Block-Elimination Algorithms for Solving Bordered Singular Systems,* Tech. Rep. 226, Computer Science Department, Yale Univ., New Haven, CT 06520, 1982. To appear in Siam J. Sci. Stat. Comp.

[4] T.F. Chan, *On the Existence and Computation of LU-factorizations with Small Pivots,* Tech. Rep. 227, Computer Science Department, Yale Univ., New Haven, CT 06520, 1982. To appear in Math. Comp.

[5] T.F. Chan, *Newton-Like Pseudo-Arclength Methods for Computing Simple Turning Points,* Tech. Rep. 233, Computer Science Dept., Yale University, New Haven, CT 06520, 1982. To appear in Siam J. Sci. Stat. Comp.

[6] G. Dahlquist and A. Bjorck, *Numerical Methods,* Prentice-Hall, Englewood Cliffs, N.J., 1974.

[7] J.J. Dongarra, J.R. Bunch, C.B. Moler and G.W. Stewart, *LINPACK User's Guide,* SIAM, Philadelphia, 1979.

[8] R.E. Funderlic and J.B. Mankin, *Solution of Homogeneous Systems of Linear Equations Arising from Compartmental Models,* SIAM J. Sci. Stat. Comp., 2 (1981), pp. 375-383.

[9] C.B. Garcia and W.I. Zangwill, *Pathways to Solutions, Fixed Points and Equilibria,* Prentice-Hall, Englewood Cliffs, N.J., 1981.

[10] P.E. Gill, W. Murray and M. Wright, *Practical Optimization,* Academic Press, New York, 1981.

[11] G.H. Golub and C. Reinsch, *Singular Value Decomposition and Least Squares Solutions,* Numer. Math., 14 (1970), pp. 403-420.

[12]
G.H. Golub and J.H. Wilkinson, *Ill-Conditioned Eigensystems and the Computation of Jordan Canonical Form,* SIAM Review, 18 (1976), pp. 578-619.

[13]
H.B.Keller, *Numerical Solution of Bifurcation and Nonlinear Eigenvalue Problems,* Applications of Bifurcation Theory, P. Rabinowitz, ed., Academic Press, New York, 1977, pp. 359-384.

[14]
H.B. Keller, Singular Systems, Inverse Iteration and Least Squares, unpublished manuscript, Applied Mathematics Department, Caltech, Pasadena, California.

[15]
H.B. Keller, Numerical Continuation Methods, Short Course Lecture Notes, National Bureau of Standards, Center for Applied Mathematics.

[16]
R.G. Melhem and W.C. Rheinboldt, *A Comparison of Methods for Determining Turning Points of Nonlinear Equations,* Computing, 29 (1982), pp. 201-226.

[17]
H.D. Mittelmann and H. Weber, *Numerical Methods for Bifurcation Problems - A Survey and Classification,* Bifurcation Problems and their Numerical Solution, Workshop on Bifurcation Problems and their Numerical Solution, January 15-17, Dortmund, 1980, pp. 1-45.

[18]
G. Peters and J.H. Wilkinson, *The Least Squares Problem and Pseudo-Inverses,* The Computer Journal, 13 (1970), pp. 309-316.

[19]
W.C. Rheinboldt, *Numerical Methods for a Class of Finite Dimensional Bifurcation Problems,* SIAM J. of Numer. Anal., 15 (1978), pp. 1-11.

[20]
W.C. Rheinboldt, *Solution Fields of Nonlinear Equations and Continuation Methods,* SIAM J. Numer. Anal., 17 (1980), pp. 221-237.

[21]
G.W. Stewart, *Computable Error Bounds for Aggregated Markov Chains,* Tech. Rep. 901, University of Maryland Computer Science Center , College Park, Maryland, 1980.

[22]
G.W. Stewart, *On the Implicit Deflation of Nearly Singular Systems of Linear Equations,* SIAM J. Sci. Stat. Comp., 2 (1981), pp. 136-140.

[23]
J.H. Wilkinson, *The Algebraic Eigenvalue Problem,* Oxford University Press, London, 1965.