

Acceleration of Convergence by Density Matrix Extrapolation in CNDO/2 Calculations

J. C. HOWLAND and R. L. FLURRY, JR.

Department of Chemistry, Louisiana State University in New Orleans,
New Orleans, Louisiana 70122, USA

Received February 16, 1972

A five-point extrapolation procedure has been applied to the first-order density matrix in CNDO/2 calculations. A number of deviations from the usual procedures are employed which protect the overall extrapolation from failure when the points used in the extrapolation of a given element form a linear or divergent series. Overall the method saves about 20–50% on the number of iterations required to attain self-consistency.

Eine Extrapolation mittels fünf Punkten wurde für die Dichtematrix 1. Ordnung bei CNDO/2-Rechnungen vorgenommen. Dabei muß man in einigen Punkten vom normalen Verfahren abweichen, um Divergenzen zu vermeiden. Man spart dann 20–50% der Iterationsschritte.

Introduction

We have been correlating the electronic structure of pharmacologically active compounds with their biological activity. This involves calculating various electronic indices within the CNDO/2 approximation [1]. In early studies we were unable to duplicate the calculated frontier electron densities and dipole moments to more than three significant figures after we had made some minor changes in the program to increase its speed and accuracy. Careful examination of the data for formaldehyde revealed that elements of the first-order density matrix were self-consistent only to within two significant figures when the electronic energy had converged to seven significant figures (Table 1). Decreasing the energy convergence criterion by a factor of ten did not significantly improve the self-consistency of the density matrix.

Since we were primarily interested in properties related to atomic charges, it was decided to converge directly on the first-order density matrix rather than on electronic energy as is done in most CNDO/2 programs. It was found that saturated molecules converged satisfactorily, but non-hydrocarbon compounds with π bonding systems took much longer to converge. In fact, 3-methylxanthane, which converged on electronic energy to a relative tolerance of 1×10^{-7} after 17 iterations failed to converge on the charge-density bond-order matrix to an absolute tolerance of 1×10^{-4} within 30 iterations. By studying the density matrix after each iteration we observed that some elements were oscillating about a value rather than converging. Therefore, if the SCF calculations were to converge within a reasonable number of iterations, some extrapolation procedure was required. Extrapolating to a "final" density matrix was chosen over extrapolating to a "final" coefficients matrix, due to the fact that the coefficients of degenerate orbitals can vary widely from iteration to iteration and still be valid solutions

Table 1. Atomic orbital charges in formaldehyde at energy convergence

Atom	Orbital	Cycle 7 ^a	Cycle 8 ^a	Diff.
H ₁	1s	1.01344	1.01344	0.00000
H ₂	1s	1.01344	1.01344	0.00000
C	2s	1.07827	1.07841	0.00014
O	2s	1.73230	1.73204	0.00026
C	2p _x	0.92842	0.92883	0.00041
O	2p _x	1.36556	1.36510	0.00046
C	2p _y	0.94167	0.94184	0.00017
O	2p _y	1.92690	1.92688	0.00002
C	2p _z	0.83472	0.83822	0.00350
O	2p _z	1.16529	1.16178	0.00351

^a At cycle 7 the electronic energy was 1.2141777×10^3 eV and at cycle 8 1.2141778×10^3 eV. The relative error was 8.23×10^{-8} .

to the Hartree-Fock equations unless symmetry is explicitly included. Elements of the density matrix do not show such variation. Small errors, representing electron loss or gain, introduced into the density matrix by the extrapolation are corrected in the next cycle of the SCF procedure. Our procedure differs from the usual procedures in three important points: a) It extrapolates on the density matrix rather than the eigenvectors. b) Extrapolation on a given elements is not carried out if that element has met the convergence criterion. c) A trap is provided to prevent extrapolation failures due to linear convergence of some elements.

Method and Application

The extrapolation procedure employed was a modification of Aitken's generalized δ^2 process as described by Winter and Dunning [2]. This procedure is a five point extrapolation which requires saving density matrices from five successive iterations.

Using the notation of Winter and Dunning, the process is described as follows:

$P_0^0 P_0^1 P_0^2 P_0^3 P_0^4$ are specific elements of the density matrix from five successive iterations. (Actually, our program uses the charge density, bond order matrix, i.e. twice the first order density matrix.) An extrapolation triangle is constructed in which successive columns of the triangle are generated by the expression

$$P_{i+1}^m = P_{i-1}^{m+1} + \frac{1}{P_i^{m+1} - P_i^m} \cdot \quad (1)$$

When $i=0$, P_{i-1}^{m+1} is set to zero. The resulting triangle is

$$\begin{array}{cccccc}
 P_0^0 & & & & & \\
 P_0^1 & P_1^0 & & & & \\
 P_0^2 & P_1^1 & P_2^0 & & & \\
 P_0^3 & P_1^2 & P_2^1 & P_3^0 & & \\
 P_0^4 & P_1^3 & P_2^2 & P_3^1 & P_4^0 & \\
 & & & & &
 \end{array} \quad (2)$$

in which P_4^0 is the final extrapolated value.

This method extrapolates successfully when the points converge asymptotically or oscillate. However, if the extrapolation is attempted on elements which have already converged, random roundoff errors frequently force extrapolation to a value outside of the convergence range. This problem was overcome by testing each element of the density matrix for convergence before extrapolating and extrapolating on only those that had not met the convergence criterion. The method also fails if the points, P_0^i , approach linearity because the denominator in Eq. (1) approaches zero at the third stage in the extrapolation. In an actual calculation, elements of the triangle are tested at different stages of the extrapolation to avoid a failure due to near linearity of the points. At the end of the second stage of the extrapolation the difference between P_2^0 and P_2^1 as well as the difference between P_2^1 and P_2^2 is compared with the convergence criterion. If either of these two differences are less than the criterion the extrapolation is stopped at this point and P_2^1 is used as the final extrapolated value. Usually when one of the differences is very small and the other difference much larger, the correction to P_2^1 to form P_4^0 , the final value, is negligible.

At the end of the third stage, P_3^0 and P_3^1 are tested to determine whether or not they are nearly equal. If the original points used in the extrapolation are nearly linear, then the difference between P_3^0 and P_3^1 will be close to zero and the extrapolation will fail. To protect against such a failure, the extrapolation is abandoned if the difference between P_3^0 and P_3^1 would yield a correction to P_2^1 that would be greater than 2 times the maximum difference between P_2^0 , P_2^1 and P_2^1 , P_2^2 , found at the end of the second stage of the extrapolation. When the extrapolation failed

Table 2. Maximum difference between elements of the current and previous density matrices and the number of elements that failed the convergence test

Formaldehyde

Cycle	Unextrapolated		Extrapolated	
	No. failures	Maximum diff. ^a	No. failures	Maximum diff. ^a
2	31	0.13016	31	0.13066
3	31	0.04955	31	0.04955
4	29	0.03931	29	0.03931
5	31	0.01130	31	0.01130
6	24	0.01147	24	0.01147
7	28	0.00313	28	0.00313
8 ^b	11	0.00350	11	0.00350
9	18	0.00111	(not tested)	
10	3	0.00115	6	0.00035
11	11	0.00041	2	0.00013
12	2	0.00040	0	0.00009
13	4	0.00015		
14	2	0.00015		
15	0	0.00006		

^a These numbers are actually for the charge-density bond-order matrix, i.e. twice the density matrix.

^b Energy converged to eight significant figures.

Table 3. Maximum difference between elements of the current and previous density matrices and the number of elements that failed to meet the convergence test

Cycle	Unextrapolated		Extrapolated	
	No. failures	Maximum diff. ^a	No. failures	Maximum diff. ^a
2	467	0.47835	467	0.47835
4	463	0.23970	463	0.23970
6	453	0.12335	453	0.12335
8	438	0.06025	438	0.06025
9	437	0.04238	(not tested)	
10	416	0.02929	215	0.00250
14	282	0.00690	41 ^b	0.00064
16	179	0.00340	18	0.00031
17	131	0.00236	(not tested)	
18	87	0.00165	0	0.00006
23	16 ^b	0.00027		
24	5	0.00019		
25	2	0.00013		
26	0	0.00009		

^a These numbers are actually for the charge-density, bond-order matrix, i.e. twice the density matrix.

^b Energy converged to eight significant figures.

at this stage, the original points were fit to a linear least squares equation (of P_0^i as a function of i) and extrapolated to one iteration past the latest iteration. Without this linearity trap, no significant improvement in convergence was obtained by use of the extrapolation. The procedure was tested with both the actual elements of five successive iterations as well as the successive averages of density matrices from six successive iterations. Averaging did not appear to improve the extrapolation. In fact, the damping due to the averaging made many of the extrapolations near enough to the linear case to yield a poor extrapolation.

Table 2 lists, for formaldehyde, the maximum difference in the element of the density matrices that was observed at each cycle and the number of elements that failed the convergence test. Since the extrapolation was done at cycle 8, no testing for convergence was done at cycle 9. Note that with the extrapolation, the problem converged in 12 iterations as opposed to 15 iterations without extrapolation. Table 3 presents similar information for cytosine.

Our subsequent experience has shown that one cycle beyond the extrapolation is sufficient for the density matrix to stabilize. Thus, the extrapolation can be carried out every six cycles. This has been chosen as our standard procedure in our CNDO/2 program. We have performed numerous calculations on systems with up to 146 basis functions. These routinely converge to 5×10^{-5} in all elements of the first-order density matrix in 20 iterations or less. For example, 3-methylxanthine converged to only 5×10^{-4} in 27 unextrapolated iterations, but to 5×10^{-5} in 20 iterations with extrapolation. Overall, the procedure saves from approximately 20 to 50% on the number of iterations required for a given convergence limit.

This extrapolation method can accelerate convergence when criteria other than electronic energy or elements of the density matrix are used. For example, calculations of water with one O-H bond two times the normal length required 34 iterations for the orbital energies to converge to an absolute error less than 1×10^{-4} . The total electronic energy converged to a relative error of less than 1×10^{-8} in 20 iterations and the elements of the density matrix converged to an absolute error less than 1×10^{-4} in 28 iterations. Extrapolating every six cycles, only 24 iterations were required for the orbital energies to converge, 14 for the total electronic energy to converge and 17 for the density matrix to converge.

This example with the water molecule also illustrates how this extrapolation procedure can accelerate convergence when the geometry is less than optimal. This procedure could prove very useful in studies designed to determine preferred conformations and structures of reaction intermediates. However, this extrapolation procedure does not help when the geometry is such that the CNDO/2 calculations would normally diverge, as, for example, in water with an O-H bond length three times normal [3]. In this instance, some diagonal elements of the density matrix maintain nearly the same absolute value, but change sign. In this case, the resulting extrapolated element would be near zero.

In the course of this study we also examined a three point procedure [4]. It was found that the three point procedure was about as good as the five point in most cases. In some instances, however, the five point procedure was clearly superior.

In order for the three-point extrapolation to be successful, the points used must be tested for linearity or divergence and suitable action taken. In our case, we fit the points to a linear least squares line and extrapolated to the next value. In the three point procedure, it was helpful to use the averages of four successive density matrices to define the three points used for the extrapolations.

Acknowledgements. The calculations were done on the Digital Equipment Corp. PDP-10 system at L.S.U.N.O. supported in part by NSF grant no. GJ-131. Partial financial support for the work came from NSF grants no. GY-2632 and GP-10418.

References

1. Pople, J.A., Beveridge, D.L.: Approximate molecular orbital theory. New York: McGraw-Hill Book Company 1970.
2. Winter, N.W., Dunning, T.H., Jr.: Chem. Physics Letters **8**, 169 (1971).
3. Jakubetz, W., Lischka, H., Rosmus, P., Schuster, P.: Chem. Physics Letters **11**, 38 (1971).
4. Roothan, C.C.J., Bagus, P.S.: Methods in computational physics, Vol. 2, p. 47. Ed. by B. Alder, S. Fernbach and M. Rotenberg. New York: Academic Press 1963.

Prof. Dr. R. L. Flurry
Department of Chemistry
Louisiana State University in New Orleans
New Orleans, Louisiana 70122, USA