Global Fringe Search Techniques for VLBI

FREDERIC R. SCHWAB AND WILLIAM D. COTTON

National Radio Astronomy Observatory*

June 1982

ABSTRACT. The usual fringe search technique for VLBI is a correlator-based method. Here, more sensitive, global techniques are presented, ones yielding simultaneous solutions for all of the delay and fringe rate parameters. The number of parameters required for delay and rate compensation for an n element array is reduced from $n^2 - n$ to $3n - 3$. Two methods are given: one is based on the Fourier transform; the other is a least-squares technique. These techniques may be viewed as variants of the self-calibration/hybrid mapping techniques which already are in widespread use.

I. INTRODUCTION

Very Long Baseline Radio Interferometry (VLBI) usually involves the use of an independent time standard at each of the widely dispersed array elements, and, ordinarily, the consequent clock error is the major source of uncertainty in determining the time at which a wavefront arrives at a given antenna. Among other sources which contribute to this uncertainty are differences in the atmospheric phase path lengths over the antennas, and errors in the assumed geometric model (i.e. errors in the baseline determination). Commonly, the cumulative error is greater than the reciprocal of the recorded bandwidth, and this error, if uncorrected, may well cause complete decorrelation of the signal. In addition, the cumulative "clock error" (which now we take to include the other sources of error mentioned above) is sufficiently variable with time that it severely limits the length of the time intervals over which coherent averaging of the data can be performed. Standard practice is to correlate the signals from each pair of antennas over a range of time lags, and then to determine the difference in the "clock errors" (delay) and the difference in the first time derivative (fringe rate) — using the data from one pair of antennas, only, for each solution for these parameters. The time intervals are chosen to be short enough that the "clock error" may be assumed approximately linear.

Here we present, as an alternative to the standard correlator-based fringe search technique, a global technique which yields a simultaneous solution for all of the delay and fringe rate parameters. For each baseline, each delay or rate parameter splits into two components, ascribable to two individual antennas. Since, in an n element interferometer array, each antenna may be a part of

*The National Radio Astronomy Observatory (NRAO) is operated by Associated Universities, Inc., under contract with the National Science Foundation.

$n-1$ interferometer pairs, a common component of error is present in the observations on the various baselines involving any given antenna. It is worthwhile, therefore, in cases of low signal-to-noise ratio (S/N), to solve simultaneously, using observations from all baselines, for these antenna-based effects. With a global technique of this sort, the total number of parameters necessary for delay and rate compensation can be reduced below the number required by the standard technique (the reduction is from $n^2 - n$ parameters to $3n - 3$). A global technique obviates the requirement that there be relatively strong fringes on every baseline to be processed. Also, the technique preserves the phase closure relations satisfied by the data. Conversely, the closure properties are disturbed any time that the standard technique is applied.

The global technique may be viewed as a variant of the self-calibration/hybrid mapping techniques (Readhead *et. al.* 1980, Cotton 1979, and Schwab 1980) which are in widespread use in aperture synthesis data reduction, inasmuch as the basic assumption on which this new method is based (see eq. [1], below) is the same. Other common features are that the global fringe search method requires an initial source model; and that, in combination with a Fourier synthesis program and a deconvolution algorithm (e.g. CLEAN), it may be applied iteratively to generate successive approximations to the radio source brightness distribution.

Following a description of the standard approach to fringe search, two global fringe search techniques are outlined below. The first global method that is described is a least-squares (LS) approach. The second, a Fourier transform (FT) method, is a straightforward generalization of the standard technique. The LS method requires good starting guesses but is more flexible than the FT method. In initial tests, a hybrid algorithm has been used, combining both of the global techniques. The FT method is used to generate starting guesses for the LS algorithm.

II. The Standard Approach

Assuming that the correlator responses have been transformed to the time and frequency domain (t, ν), and that n antennas comprise the interferometer array, the standard approach to fringe search can be described as follows. Visibility observations $\tilde{V}_{ij}(t_k, \nu_l)$ on the i–j baseline $(1 \leq i < j \leq n)$ at times t_k, $k = 0, \ldots, n_t$, and at frequencies ν_l, $l = 0, \ldots, n_\nu$, are related to the true source visibility $\mathcal{V}_{ij}(t, \nu)$ according to

$$\tilde{V}_{ij}(t_k, \nu_l) = g_i(t_k, \nu_l)\overline{g}_j(t_k, \nu_l)\mathcal{V}_{ij}(t_k, \nu_l) + \epsilon_{ijkl} \qquad (1)$$

where the antenna effects (clock errors, atmospheric effects, etc.) have been absorbed into complex-valued functions $g_q(t, \nu)$, and where the observational errors ϵ are sufficiently small and well-behaved. Each g can be written as $g_q(t, \nu) = a_q(t, \nu)\exp[i\psi_q(t, \nu)]$ (ψ_q is called the *antenna phase* of antenna q). Equation (1) is the model relation which is exploited by the usual self-calibration/hybrid

mapping techniques. (Properly speaking, the $\tilde{V}_{ij}(t_k, \nu_l)$ are not really samples of the first term in the right-hand side of eq. [1]; rather, they are averages of this term, over some neighborhood of (t_k, ν_l). This fact may be ignored for purposes of the present discussion.) Thompson (1980) and Thompson and D'Addario (1982) present rigorous analyses of the effect of hardware electronic design considerations upon the validity of this model; further analysis is given in (Clark 1981). The so-called "phase closure" and "amplitude closure" relations follow by taking logarithms of ratios of equation (1), substituting into the relation antenna indices chosen to define a closed path among the antennas.

Although the a_q are functions of time and frequency, it is assumed that they change slowly enough that $g_q(t, \nu) = a_q \exp[i\psi_q(t, \nu)]$, with a_q constant over the (t, ν)-interval of the fringe search. Also, $|\mathcal{V}_{ij}|$ is assumed to be constant over the relevant interval. Then, to first-order,

$$\begin{aligned}\tilde{V}_{ij}(t_k, \nu_l) \simeq\ & a_i a_j \mathcal{V}_{ij}(t_0, \nu_0) \exp\{i[(\psi_i - \psi_j)(t_0, \nu_0)]\} \\ & \exp\left\{ i \left[\left. \frac{\partial(\psi_i - \psi_j + \phi_{ij})}{\partial t} \right|_{(t_0, \nu_0)} (t_k - t_0) \right.\right. \\ & \left.\left. + \left. \frac{\partial(\psi_i - \psi_j + \phi_{ij})}{\partial \nu} \right|_{(t_0, \nu_0)} (\nu_l - \nu_0) \right] \right\},\end{aligned} \tag{2}$$

where $\phi_{ij} \equiv \arg \mathcal{V}_{ij}$.

The quantities

$$r_{ij} \equiv \left. \frac{\partial(\psi_i - \psi_j + \phi_{ij})}{\partial t} \right|_{(t_0, \nu_0)} \quad \text{and} \quad \tau_{ij} \equiv \left. \frac{\partial(\psi_i - \psi_j + \phi_{ij})}{\partial \nu} \right|_{(t_0, \nu_0)}, \tag{3a, b}$$

which are called, respectively, the *fringe rate* and the *delay* for baseline i–j at (t_0, ν_0), can be estimated by searching for the location of the maximum modulus of the Fourier transform $\hat{F}_{ij}(r, \tau)$ of the distribution $F_{ij}(t, \nu)$ defined by

$$F_{ij}(t, \nu) = \sum_{\substack{0 \le k \le n_t \\ 0 \le l \le n_\nu}} \delta[t - (t_k - t_0), \nu - (\nu_l - \nu_0)] \tilde{V}_{ij}(t_k, \nu_l), \tag{4}$$

where δ is the two-dimensional Dirac delta function (i.e., the distribution with the property that $\hat{\delta}(r, \tau) \equiv 1$).* Furthermore,

$$\hat{F}_{ij}(r_{ij}, \tau_{ij}) \simeq a_i a_j \mathcal{V}_{ij}(t_0, \nu_0) \exp\{i[(\psi_i - \psi_j)(t_0, \nu_0)]\}. \tag{5}$$

The quantities r_{ij} and τ_{ij} (actually, estimates thereof), then, are the product of the standard fringe search technique. They allow one to produce phase-corrected

*In practice, one doesn't compute $\hat{F}_{ij}$ directly, but rather a discrete approximation to it—via the FFT algorithm—, and then interpolation yields the parameter estimates.

observations

$$\tilde{V}_{ij}(t_k, \nu_l)\exp\{-i[(t_k - t_0)r_{ij} + (\nu_l - \nu_0)\tau_{ij}]\} \tag{6}$$

which, to the extent that the first-order model (eq. [2]) is valid — *and* that ϕ_{ij} is nearly constant —, can be averaged coherently over time and frequency. Note carefully, though, that the technique does not separate the $\partial\psi/\partial t$ and $\partial\phi/\partial t$ components of the rates, nor the $\partial\psi/\partial\nu$ and $\partial\phi/\partial\nu$ components of the delays, as one might wish.

Following the fringe search, the corrupting influence of the quantities $a_q^{-1}\exp[-i\psi_q(t_0, \nu_0)] = g_q^{-1}(t_0, \nu_0)$ is removed in post-processing by the self-calibration/hybrid mapping techniques.

The major disadvantage of the standard approach is that relatively high S/N (i.e., small ϵ) is required on each baseline in order to obtain reliable estimates of all the r_{ij} and τ_{ij}. The obvious approach toward a fringe processing method for the low S/N regime is to try to separate the antenna-based components $\frac{\partial\psi_q}{\partial t}|_{(t_0,\nu_0)}$ and $\frac{\partial\psi_q}{\partial\nu}|_{(t_0,\nu_0)}$ of the fringe rates and the delays by means of some simultaneous solution over all baselines. Obviously, the source phase ϕ_{ij} cannot be separated into antenna-based components — so a source model V_{ij}, approximating $\mathcal{V}_{ij}$, is required, *a la* self-calibration/hybrid mapping, in the two global fringe search methods described below.

III. A Least-Squares Approach

Here, given a source model V approximating $\mathcal{V}$, above, we consider the problem of least-squares estimation of the antenna-based components of fringe rate and delay. Simultaneously we wish to estimate the $g_q(t_0, \nu_0)$. The antenna-based parameters are the a_q, the antenna phases $\psi_{q0} \equiv \psi_q(t_0, \nu_0)$, the antenna rates $r_q \equiv \frac{\partial\psi_q}{\partial t}|_{(t_0,\nu_0)}$, and the antenna delays $\tau_q \equiv \frac{\partial\psi_q}{\partial\nu}|_{(t_0,\nu_0)}$. Define

$$E_{ijkl} = \exp\{i[(\psi_{i0} - \psi_{j0}) + (r_i - r_j)(t_k - t_0) + (\tau_i - \tau_j)(\nu_l - \nu_0)]\}. \tag{7}$$

We shall fit the ensemble of observations in the (t, ν)-sample space to the model

$$\tilde{V}_{ij}(t, \nu) = a_i a_j V_{ij}(t, \nu)E_{ijkl} + \text{noise}. \tag{8}$$

Since only pairwise differences of the ψ_{q0}, the r_q, and the τ_q appear in the model, we shall settle, in the absence of other assumptions, upon estimating $4n - 3$ parameters, omitting ψ_{p0}, r_p, and τ_p for some choice p of "reference antenna." Thus, if we set $\psi_{p0} = r_p = \tau_p = 0$, then we estimate antenna phase, fringe rate, and delay with respect to the phase, rate, and delay at antenna p. A least-squares formulation of the problem, then, is to minimize the functional

$$S(\mathbf{x}) = \sum_{\substack{0\le k\le n_t \\ 0\le l\le n_\nu}} \sum_{1\le i<j\le n} \frac{1}{\sigma_{ijkl}^2}|\tilde{V}_{ij}(t_k, \nu_l) - a_i a_j V_{ij}(t_k, \nu_l)E_{ijkl}|^2. \tag{9}$$

Here, $\mathbf{x}$ is the (column) vector of unknown parameters, $\mathbf{x} = \text{col}'(a_1, \ldots, a_n, \psi_{10}, \ldots, \psi_{n0}, r_1, \ldots, r_n, \tau_1, \ldots, \tau_n)$, the prime denoting that antenna p's parameters (except for a_p) are omitted; and σ^2_{ijkl} is an estimate of the variance of $\tilde{V}_{ij}(t_k, \nu_l)$.

One might choose to minimize, instead of S, the functional

$$S_1(\mathbf{x}) = \sum_{k,l} \sum_{i<j} \frac{1}{\sigma^2_{ijkl}} |\tilde{V}_{ij}(t_k, \nu_l)/V_{ij}(t_k, \nu_l) - a_i a_j E_{ijkl}|^2 \tag{10}$$

— with an appropriate redefinition of the weights (this form may be preferred when data storage is at a premium); or, if it were assumed that $a_q = 1$ for every q,

$$S_2(\mathbf{x}) = \sum_{k,l} \sum_{i<j} \frac{1}{\sigma^2_{ijkl}} |\tilde{V}_{ij}(t_k, \nu_l)/V_{ij}(t_k, \nu_l) - E_{ijkl}|^2, \tag{11}$$

where the a_q have been deleted from the vector $\mathbf{x}$ of unknowns; or, if there were a desire to neglect the amplitude information

$$S_3(\mathbf{x}) = \sum_{k,l} \sum_{i<j} \frac{1}{\sigma^2_{ijkl}} |\exp\{i[\tilde{\phi}_{ij}(t_k, \nu_l) - \phi_{ij}(t_k, \nu_l)]\} - E_{ijkl}|^2, \tag{12}$$

where $\tilde{\phi}_{ij} \equiv \arg \tilde{V}_{ij}$, and where, again, the a_q have been deleted from $\mathbf{x}$. Thus far, in the attempt at a practical implementation of the LS method, only the form S_3 has been employed.

Additionally, one might choose to incorporate prior knowledge of the parameters: for example, one might tabulate a running mean m_{r_q} of solutions for each antenna's fringe rate r_q, along with an empirical measure of the r.m.s. scatter σ_{r_q} about this running mean. If one were then to assume a Gaussian prior distribution $\mathcal{N}(m_{r_q}, \sigma^2_{r_q})$ for r_q, one would add to S the "penalty" term

$$\sum_{\substack{q=1 \\ q \neq p}}^{n} \frac{(r_q - m_{r_q})^2}{\sigma^2_{r_q}}. \tag{13}$$

Similar terms would be added for the other parameters. (Cornwell and Wilkinson (1981) advocate just such a Bayesian approach in ordinary self-calibration/hybrid mapping). Or, if there were simply a desire to constrain the estimates of some parameter to a given interval, one could add to S a differentiable penalty function which is small over the given interval and which rises sharply at its boundaries. Either of these techniques would serve to reduce the variance in the parameter estimates — at the expense of increased bias, over and above the bias inherent in the basic nonlinear LS approach. Preliminary tests of the method indicate that constraints are not required in cases of moderate S/N in order for the method to behave sensibly; but, since often there is good prior knowledge, the present

implementation of the method does allow the use of prior distributions for the r_q and the τ_q.

S typically has many local minima, as well as other critical points. One can reasonably expect that a standard iterative minimization algorithm, starting with a reasonable initial guess, reliably and efficiently should be able to locate some local minimum. Locating a starting point for which a *global* minimum will be the point of attraction of the algorithm is more difficult. However, secondary minima of S ought to be spaced in fringe rate and delay roughly like the maxima of the cosine transforms of $f(t) = \sum_k w^\dagger(k)\delta(t - t_k + t_0)$ and $g(\nu) = \sum_l w^\ddagger(l)\delta(\nu - \nu_l + \nu_0)$ whenever the variance estimates are separable functions of k and l; i.e., whenever $\sigma_{ijkl} = \sigma_{ij}/\sqrt{w^\dagger(k)w^\ddagger(l)}$. In tests with model data, this observation has improved the success rate of the algorithm substantially when started with arbitrary initial guesses. But a more fruitful approach is to use a generalization of the standard method of § II in order to generate the starting guesses. This scheme is described in the next section.

A detailed description of an algorithm which has proven effective in solving the LS problem, as formulated above, is given in the Appendix.

IV. A Generalization of the Standard Approach

As in the least-squares minimization of S_3, let us discard the amplitude information but retain the phases, $\tilde{\phi}$ and ϕ. Then, to estimate the ψ_{q0}, the r_q, and the τ_q (but not the a_q), a global Fourier transform approach analogous to the standard approach is easily derived.

For simplicity of exposition, assume that antenna number 1 has been chosen as reference antenna; i.e., $p = 1$. Define the distributions

$$D_{kl}(t, \nu) = \delta[t - (t_k - t_0), \nu - (\nu_l - \nu_0)]. \tag{14}$$

Then the location of the maximum modulus of the Fourier transform of the distribution

$$F_{12}(t, \nu) = \sum_{k,l} \frac{1}{\sigma^2_{12kl}} D_{kl}(t, \nu) \exp\{i[(\tilde{\phi}_{12} - \phi_{12})(t_k, \nu_l)]\} \tag{15}$$

provides an estimate of τ_2 and τ_2, and the argument there is an estimate of ψ_{20}. By the phase closure relations (cf. Readhead *et. al.* 1980), $\tilde{\phi}_{12} - \phi_{12} = \tilde{\phi}_{13} - \phi_{13} - \tilde{\phi}_{23} + \phi_{23} + \text{noise}$. In general, for $j \neq 1, 2$, define

$$F_{1j2}(t, \nu) = \sum_{k,l} \frac{1}{\sigma^2_{1jkl} + \sigma^2_{2jkl}} D_{kl}(t, \nu) \exp\{i[(\tilde{\phi}_{1j} - \phi_{1j} - \tilde{\phi}_{2j} + \phi_{2j})(t_k, \nu_l)]\}. \tag{16}$$

Since the differences $\tilde{\phi} - \phi$ occurring in F_{12} and in the F_{1j2} are all independent (apart from whatever scheme was used to derive the model) the Fourier transform

of $F \equiv F_{12} + \sum_{j \neq 1,2} F_{1j2}$ yields, in general, better parameter estimates than any one of the $F_{\_}$'s alone.

Paths like 1–3–4–2, with functions F_{1jk2} defined analogously to those above, provide further estimates of τ_2, τ_2, and ψ_{20}. F_{1342} is not independent of F_{132} and F_{142}, although it incorporates new data from the 3–4 baseline. So it is not clear how to sum all of the $F_{\_}$'s together with properly chosen weights. Nevertheless, one can adopt an *ad hoc* weighting scheme. For an arbitrary choice p of reference antenna, the paths to antenna l which incorporate at least partially independent phase differences can be enumerated as:

$$(p\text{–}l); \qquad (p\text{–}j\text{–}l), \quad j = 1, \ldots, n, \quad j \neq p, l;$$

and $\quad (p\text{–}j\text{–}k\text{–}l), \quad j = 1, \ldots, n-1, \quad k = j+1, \ldots, n, \quad j, k \neq p, l.$

As in the standard fringe search technique, one would, in implementing this method, compute a discrete approximation to $\hat{F}$ — via the FFT algorithm —, rather than $\hat{F}$ itself. Given prior knowledge of the fringe rates and the delays, it would be straightforward to constrain the parameter estimates by doing a restricted search for the maximum modulus of $\hat{F}$. It is not readily apparent how the method could be generalized to provide estimates of the a_q.

The global FT algorithm is well-suited to implementation in a high-speed array processor.

V. Discussion

A hybrid global fringe search algorithm has been implemented within the AIPS interactive data reduction system. The FT method is used to provide starting guesses for the LS algorithm. In the LS routine, the functional S_3 is minimized. The weights $1/\sigma^2_{ijkl}$ are chosen to be proportional to the product $w_i w_j$ of two "antenna weights," times the square of the modulus of the model visibility, $|V_{ijkl}|^2$. Thus, baselines on which the source is heavily resolved are given small weight. A data editing capability based on goodness-of-fit to the model (eq. [8]) also has been incorporated. Typically, the FT method provides good enough parameter estimates for the LS algorithm to converge in 4-6 iterations. *A priori* knowledge of the fringe rates and the delays often is good enough that there is no difficulty in locating a global minimum of S.

As an alternative to the LS algorithm, one could use instead an ℓ_1 minimization method, in order to reduce the sensitivity to wild data points. This approach, whose application to ordinary self-calibration/hybrid mapping is described in (Schwab 1981), readily could be adapted to the problem at hand.

The global fringe fitting technique offers a number of substantial advantages over the standard baseline oriented fringe fitting method. Among them are:

1) By its very nature, the new method forces closure of the derived instrumental delays and rates, whereas the standard technique, in general, does not preserve delay and rate closure.

2) When, as is common, there are great differences in the sensitivity of the antennas comprising the VLBI array, the global technique allows the use of data

on baselines which by themselves are much too insensitive for baseline fringe fitting to succeed.

3) For any given array, the global technique lowers the threshold for minimum detectable point source flux density.

4) Global fringe fitting offers a simple way to process polarization observations, as the delays and rates for each I.F. in use can be determined from the parallel-polarized correlations, and then they can be applied to the cross-hand polarized correlations. This removes the requirement that the source be detectable in the cross-hand polarized channels.

5) The global fringe fitting technique allows the mapping of sources whose size is not small compared to the delay resolution on every baseline. This follows because the source structure information can be used in an iterative procedure to remove the effects of source structure before fitting for antenna delays and rates. The source then can be mapped in each frequency channel separately, and the resulting maps averaged. The larger delay beam of the narrow band frequency channels becomes the limitation, rather than the delay beam corresponding to the full bandpass. Since the frequency resolution can be set arbitrarily by changing the numbers of delays correlated, the effective delay beam can be made as large as desired.

6) The increased sensitivity of the global fringe fitting technique allows the use of shorter solution intervals than can be used for baseline fringe fitting. Hence, the actual variations of delay and rate are better approximated by the linear model (eq. [2]), than they are in the standard technique. In most cases this results in more accurate phase correction of the measured correlations.

APPENDIX

In this section, S denotes whichever functional of § III, S, S_1, S_2, or S_3, is to be minimized. $\mathbf{x}$ is the corresponding column vector of unknown parameters. With n antennas, there are $N = 4n - 3$ parameters unless S_2 or S_3 has been selected, in which case $N = 3n - 3$. Let $\nabla S(\mathbf{x})$ denote the gradient of S evaluated at $\mathbf{x}$; i.e., $\nabla S(\mathbf{x}) = \text{col}\left(\frac{\partial S}{\partial x_1}, \ldots, \frac{\partial S}{\partial x_N}\right)$; and let $H(\mathbf{x})$ denote the Hessian matrix of S at $\mathbf{x}$; i.e., $H(\mathbf{x}) = \left(\frac{\partial^2 S}{\partial x_i \partial x_j}\right)_{i,j=1}^{N}$. In general, S has many critical points — local minima, local maxima, and saddle points. At any critical point $\mathbf{x}^*$, $\nabla S(\mathbf{x}^*) = 0$. At a local minimum, $H(\mathbf{x}^*)$ is positive semidefinite (i.e., all of its eigenvalues are nonnegative), and at a proper local minimum (i.e., when all parameters are well-determined), $H(\mathbf{x}^*)$ is positive definite. At a local maximum, H is negative semidefinite (i.e., none of its eigenvalues is positive), and at a saddle point H is indefinite (i.e., it has some negative and some positive eigenvalues). Away from a critical point, $-\nabla S(\mathbf{x})$ points in the direction of steepest descent in the sense of the standard Euclidean metric $\sqrt{(\mathbf{x}-\mathbf{y})^{\mathrm{T}}(\mathbf{x}-\mathbf{y})}$. For any positive definite $N \times N$ matrix G, $-G\nabla S(\mathbf{x})$ is also a steepest descent direction — steepest in the sense of the metric $\sqrt{(\mathbf{x}-\mathbf{y})^{\mathrm{T}} G (\mathbf{x}-\mathbf{y})}$ derived from the inner product $\mathbf{x}^{\mathrm{T}} G \mathbf{y}$.

Clearly, because of the multitude of local maxima and saddle points, we require a descent algorithm. One would not generally choose, though, to step in the direction of steepest descent in the sense of the Euclidean metric because such a strategy can lead to arbitrarily slow convergence. But whenever $H(\mathbf{x})$ is positive definite, $-H^{-1}(\mathbf{x})\nabla S(\mathbf{x})$ is a steepest descent direction. And steps taken in this search direction will converge eventually at a *quadratic* rate in the neighborhood of any proper local minimum of S (cf. Ortega and Rheinboldt (1970) or Murray (1972)).

Because of the special structure of our problem, a simple variant of Newton's method for root-finding, applied to the equation $\nabla S(\mathbf{x}) = 0$, is practical. A variant is required because we want the algorithm to converge upon only one type of critical point — a local minimum. The reason that the structure is special and that Newton's method is practical is that, when $N = 3n - 3$, each observation, since it involves only two antennas, and, therefore, at most six unknown parameters, contributes to at most six elements of ∇S, to at most six of the diagonal elements of H, and to at most 15 subdiagonal elements of the (symmetric) matrix H. Furthermore, all of these contributions are easy to compute. When $N = 4n - 3$, the contributions are to at most 8, 8, and 28 elements, respectively. For many other problems, Newton's method is impractical because each observation may involve all of the parameters and because analytic second order derivatives of the model may be very difficult to obtain.

The raw form of Newton's algorithm is: Given an initial guess $\mathbf{x}^{(0)}$, at the k^{th} iteration form the new parameter estimate $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \alpha_k H^{-1}(\mathbf{x}^{(k)})\nabla S(\mathbf{x}^{(k)})$, where $\alpha_k > 0$ is chosen so that $S(\mathbf{x}^{(k+1)}) \leq S(\mathbf{x}^{(k)})$. The difficulty is that H^{-1} (and equivalently H) must be positive definite in order for $-H^{-1}\nabla S$ to be a descent direction. Hence, any time that $H(\mathbf{x}^{(k)})$ has nonpositive eigenvalues, we shall wish to replace it by $\tilde{H}(\mathbf{x}^{(k)})$, where $\tilde{H}$ is positive definite, and in some sense close to H. One means of doing so, known as Greenstadt's modification (Murray 1975, p.59), is to compute the eigenvector expansion of H, $H = \sum_{i=1}^{N} \lambda_i \mathbf{v}_i \mathbf{v}_i^{\mathrm{T}}$, with $\mathbf{v}_i$ orthonormal, and modify it slightly so that, for some small $\epsilon > 0$,

$$\tilde{H} = \sum_{i=1}^{N} \tilde{\lambda}_i \mathbf{v}_i \mathbf{v}_i^{\mathrm{T}}, \quad \text{where } \tilde{\lambda}_i \equiv \begin{cases} \lambda_i & \text{if } \lambda_i \geq \epsilon \\ \epsilon & \text{otherwise.} \end{cases} \tag{17}$$

(On a t digit base β machine, $\epsilon \approx \beta^{-t}$ is an appropriate choice, since positive eigenvalues smaller than this cannot easily be distinguished from other eigenvalues that are close to zero). Then the search direction is given by

$$-\tilde{H}^{-1}\nabla S = -\sum_{i=1}^{N} \frac{1}{\tilde{\lambda}_i}[\mathbf{v}_i^{\mathrm{T}}\nabla S]\mathbf{v}_i. \tag{18}$$

Note that $\mathbf{v}_i^T \nabla S$ is a scalar — the dot product of the i^{th} normalized eigenvector of H with ∇S.

A rough block diagram of an implementation of the LS algorithm is shown in Figure 1. The subroutine DSICO of the LINPACK package (Dongarra *et. al.* 1979) is used to factor H, $H = UDU^T$, with U unit upper triangular and D a diagonal matrix. Using this factorization, DSIDI, also taken from LINPACK, tallies the number of positive, negative and zero eigenvalues of H. When H is positive definite, the subroutine DSISL of LINPACK is used to compute the search direction $\mathbf{d}_k = -H^{-1}(\mathbf{x}^{(k)}) \nabla S(\mathbf{x}^{(k)})$.

Whenever DSIDI reports that there are nonpositive eigenvalues, Greenstadt's modification is applied. The subroutine TRED2 of the EISPACK package (Smith *et. al.* 1976) is used to reduce H to tridiagonal form, via similarity transforms. Then IMTQL2, also of EISPACK, is used to compute the eigensystem by the implicit QL method. The implicit QL algorithm is quite forgiving of bad parameter scaling — say, if the parameters are defined in units in which the fringe rates are orders of magnitude different from the antenna phases — this can lead to eigenvalues of widely varying magnitude. Once the eigensystem has been computed, the search direction is given by $\mathbf{d}_k = -\tilde{H}^{-1}(\mathbf{x}^{(k)}) \nabla S(\mathbf{x}^{(k)})$.

Finally, a coarse search is made in the direction $\mathbf{d}_k$, for an α_k satisfying $S(\mathbf{x}^{(k)} + \alpha_k \mathbf{d}_k) \leq S(\mathbf{x}^{(k)})$. Close enough to a proper local minimum, the choice $\alpha_k = 1$ will do. We set $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha_k \mathbf{d}_k$, test for convergence, and perform another iteration, if necessary. The theoretically predicted quadratic rate of convergence near a proper local minimum frequently has been observed in the initial tests of the algorithm.

Because the first and second derivatives of the model function are entirely straightforward to compute, the expressions for ∇S and for H have been omitted here.

Preliminary results have shown this method to perform well, but other implementations of the LS algorithm would do as well or better. In particular, an alternative method of modifying indefinite H could speed up the algorithm significantly, since the eigen-decomposition is relatively expensive. The algorithm given here could be implemented in a high-speed array processor, at a moderate expenditure of human effort.

REFERENCES

Clark, B. G. (October 1981), VLA Scientific Memo. No. 137 (VLA, P.O. Box O, Socorro, NM 87801).

Cornwell, T. J. and Wilkinson, P. N. (1981), *Mon. Not. R. Astron. Soc.*, **196**, 1067.

Cotton, W. D. (1979), *Astron. J.*, **84**, 1122.

Dongarra, J. J., Bunch, J. R., Moler, C. B., and Stewart, G. W. (1979), **LINPACK Users' Guide**, SIAM, Philadelphia.

Murray, W. (1972), **Numerical Methods for Unconstrained Optimization**, Acad. Pr., New York.

Ortega, J. M. and Rheinboldt, W. C. (1970), **Iterative Solution of Nonlinear Equations in Several Variables**, Acad. Pr., New York.

Readhead, A. C. S., Walker, R. C., Pearson, T. J., and Cohen, M. H. (1980), *Nature*, **285**, 137.

Schwab, F. R. (1980), In W. T. Rhodes (ed.), "1980 International Optical Computing Conference", *Proc. Soc. Photo-Opt. Instr. Eng.*, **231**, pp. 18–25.

Schwab, F. R. (September 1981), VLA Scientific Memo. No. 136 (VLA, P.O. Box O, Socorro, NM 87801).

Smith, B. T., Boyle, J. M., Dongarra, J. J., Garbow, B. S., Ikebe, Y., Klema, V. C., and Moler, C. B. (1976), **Matrix Eigensystem Routines — EISPACK Guide**, Lecture Notes in Computer Science 6, Springer-Verlag, New York.

Thompson, A. R. (February 1980), VLA Electronics Memo. No. 192 (VLA, P.O. Box O, Socorro, NM 87801); Addendum, (March 1980).

Thompson, A. R. and D'Addario, L. R. (1982), *Radio Science*, **17**, 357.

National Radio Astronomy Observatory, Edgemont Road, Charlottesville, VA 22901

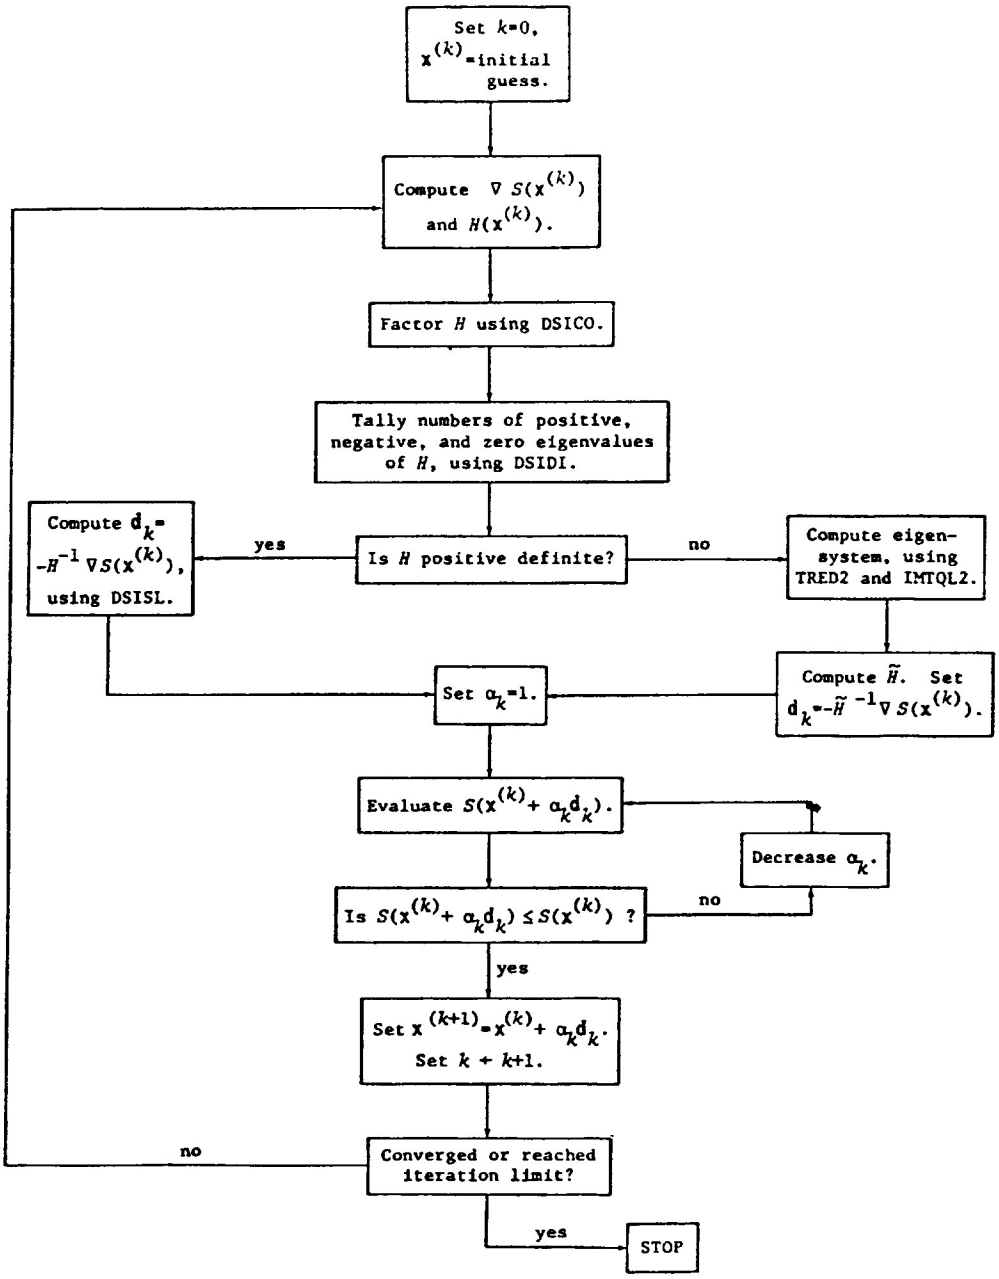

Figure 1. Rough block diagram of the LS algorithm.