Peak Tracking by Simultaneous Inversion: Toward a One-Step Acoustic Tomography Analysis

UWE SEND

Institut für Meereskunde an der Universität Kiel, Kiel, Germany

4 May 1995 and 18 January 1996

ABSTRACT

A number of geophysical observing techniques, including ocean acoustic tomography, obtain sequences of records of which the observed relative maxima ("peaks") are used to infer properties of the system via inversions. Traditionally, these peaks first are tracked (followed from one record to another) and identified separately, before they can be used in an inversion scheme. In this paper, a method is presented for identifying and tracking ensembles of such peaks in one step and simultaneously with the inversion. A priori information, in our case knowledge about the ocean, can thus be used to constrain the allowed peak identifications, enabling the usage of irregularly appearing or more closely spaced peaks. The best identification is defined to be the one that upon inversion minimizes a cost function that involves data residual and smoothness in time, subject to two constraints bounding the solution and residual size. For the presented cases, the minimum can be found by simply trying inversions with *all* possible peak identifications. Sample applications of the method from an acoustic tomography experiment are shown in order to illustrate the approach and results.

1. Introduction

Consider a series of geophysical records, each yielding a set of relative maxima, called "peaks," whose properties (amplitude, location) constitute the basic data. An example is shown in Fig. 1, taken from the ocean acoustic tomography (e.g., Munk and Worcester 1988) experiment that yielded the data used in this paper. In such techniques, the peaks usually need to be identified before they can be used for inferring the properties of the system that is being investigated. This is often done by a tracking algorithm, which seeks to follow a certain peak from one sample to another, assuming some smooth variation. Such peak tracking has, at least in acoustic tomography, always been a cumbersome and time-consuming process. In addition, it is artificially detached from the actual problem of inverting the data.

In this study, a more general method is presented. It is particularly useful when tracking becomes ambiguous, for example, due to peaks fading in and out, combined with the theoretically possible peaks becoming as close as the resolution limit. A tracking algorithm that only tries to recognize the pattern or searches for the maximum inside a window around the last peak location will not be able to follow a certain peak if it is surrounded closely by many others and if it (or the neighbors) in addition is not always visible. However, one usually has some a priori knowledge about the system affecting the peaks (here the sound speed structure of the ocean), which constrains the allowed arrangement relative to one another. This knowledge can be exploited to try to make objective choices during peak tracking or identification. Since this reduces to a minimization problem, the peak identification can in fact build upon the inversion of the peak data and be carried out simultaneously with it. Information about the time evolution of the peaks can also be included in the minimization.

2. Principle of algorithm

In the following it will be assumed that (as in acoustic tomography) it is the *location* (here travel time) of the peaks that forms the basic data. No amplitude or full-wave inversions will be considered. Further, different terms will be used to make clear the distinction between actually observed *peaks* and the expected or predicted possible *arrivals*. Ideally, the observed peaks will be identical with the expected arrivals. If there is a different number of observed peaks than predicted arrivals (or if their order is not unique), each observed peak needs to be identified as one of the expected arrivals. The linearized theoretical relation between the system and the data can be written as

$$\mathbf{d} = \mathbf{G}\mathbf{m} + \mathbf{e}. \quad (1)$$

Here the vector $\mathbf{m}$ describes the system properties sought and the observation matrix $\mathbf{G}$ gives the depen-

Corresponding author address: Dr. Uwe Send, Institut für Meereskunde an der Universität Kiel, Abt. Reg. Ozeanographie, Düsternbrooker Weg 20, D-24105 Kiel, Germany.
E-mail: usend@ifm.uni-kiel.de

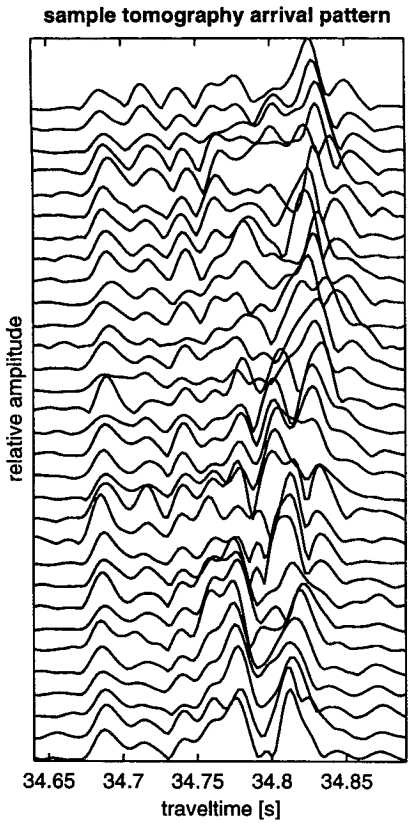

FIG. 1. One-day sequence of peak arrivals from the ocean acoustic tomography experiment used in this paper for a sample application. Each reception is offset vertically by a fixed amount. The peaks correspond to discrete acoustic pulses (from different acoustic ray paths) whose travel times yield a measurement of the sound speed structure between transmitter and receiver.

dence of the observed data $\mathbf{d}$ on $\mathbf{m}$. The elements d_i of the data vector are the locations of I theoretically possible arrivals i, of which we assume a total of J to have been observed as peaks. The measurement errors are given by the elements e_i of the vector $\mathbf{e}$. When a set of observed travel times d_j is available—that is, after peak identification—standard inverse techniques (e.g., Aki and Richards 1980) can be applied to obtain an estimate $\hat{\mathbf{m}}$ for the system parameters. The inversion can be understood as seeking a solution that minimizes the data misfit (residuals) and the solution variance. Usually, we build a lot of our knowledge about the system into the solution, for example, by using a description $\mathbf{m}$ in terms of dynamical or statistical modes resulting in a reduced number of degrees of freedom and by specifying the covariance of $\mathbf{m}$, as well as that of the data errors $\mathbf{e}$ (e.g., Howe et al. 1987).

Now consider a measurement of J peaks with travel times d_j, where it is not clear which of the expected arrivals i these correspond to. The basic assumption of the method is that usually *all* observed peaks of significant height can be identified with a theoretically possible arrival (of ray, diffracted, or modal nature), but that the subset of possible arrivals that is visible at any one time may change.[1] For the identification we can make use of our knowledge about the system. For example, the known covariance of $\mathbf{m}$ can be translated into a covariance of the arrival locations $\mathbf{d}$, which thus describes the dependence of arrival locations on one another. Instead of using this for a probabilistic solution based on probability density functions for peak locations or spacings, however, a simpler approach is adopted here. It exploits the expectation that "unlikely" peak arrangements would yield equally unlikely (large) solutions $\hat{\mathbf{m}}$, if an inversion were tried. Thus, two conditions are set for the optimum peak identification: Upon performing an inversion for a particular trial identification, 1) the resulting solution $\hat{\mathbf{m}}$ must be within the realistic range of amplitudes (two to three standard deviations or absolute bounds); 2) the resulting data residuals $\mathbf{d} - \hat{\mathbf{d}}$, that is, the misfits between the observed data and the ones predicted from the inverse solution must be within their expected bounds.

Subject to these two constraints, the best identification is chosen by optimizing a cost function that has two terms. One term builds upon the minimization implicit in the inversion that is carried out with an identification; that is, we minimize the data residual not only to find the "best" solution for $\mathbf{m}$ but also to find the best peak identification. The second term requires a smooth evolution in time of the system, expressed here in terms of the *solution* (as opposed to using peak positions, the smooth evolution of which forms the basic principle of the pattern tracking methods).

The above peak identification process is a highly nonlinear procedure in terms of the minimizations required. For the applications in this study, it turns out to be feasible to try *all* reasonable possibilities of peak identification and to choose the one that optimizes the cost function subject to the two constraints.

3. Details of method

To initialize a tracking run, a first-guess $\mathbf{m}^0$ of the mode amplitudes must be available. With that, the forward problem (1) is solved to predict the travel times of all expected arrivals d_i^0. In the first observation, a large window is set around each observed peak, and all predicted arrivals found within each window are considered possible identifications of the corresponding observed peak. The result is a table, which for each observed peak has a list of possible identifications.

[1] There may not be a one-to-one correspondence between measured peaks and deterministic ray arrivals for propagation over very long ranges, due to internal wave scattering (Colosi et al. 1995). In such cases the hope would be to recover the some of the expected deterministic arrivals by averaging over a sufficient number of measurements.

From this table, a recursive routine then generates all possible permutations to identify the peaks. Note that the aim is to identify all observed peaks that have predicted arrivals in their window. If there is an inconsistency in this approach, for example, if only ray arrivals have been calculated but a diffracted peak is observed, the algorithm should normally reject the peak.

An example might clarify the process. Assuming that six arrivals (numbered $1, \ldots, 6$) are theoretically possible, and three peaks (labeled a, b, c) have been observed in a particular realization, the window selection might yield the following possibilities for the three observed peaks (see Table 1). Observed peak a could be expected arrival 1, peak b could be arrival 3 or 4, and peak c could be arrival 4, 5, or 6 (zero is a fill character). The possible permutations for identifying the three observed peaks thus are arrivals $\{1, 3, 4\}$, $\{1, 3, 5\}$, $\{1, 3, 6\}$, $\{1, 4, 5\}$, $\{1, 4, 6\}$. An inversion will be tried for these five possible identifications to determine the estimated mode amplitudes $\hat{\mathbf{m}}$ each time. Then the choice is made to determine the best of the peak identifications. To start with, only those identifications are retained that satisfy the two constraints, that is, yielding solution mode amplitudes $\hat{\mathbf{m}}$ within two to three standard deviations of their expected variability and data residuals $\mathbf{d} - \hat{\mathbf{d}}$ within their expected bounds.

For the remaining possibilities, the one is chosen that minimizes the cost function of the form

$$C = \frac{1}{J}\sum_{j=1}^{J}\left(\frac{\hat{d}_j - d_j}{e_j}\right)^2 + \frac{\alpha^2}{L}\sum_{l=1}^{L}(\hat{m}_l - m_l^{\text{last}})^2. \quad (2)$$

The first sum corresponds to the mean squared data residuals weighted by the inverse data error $\mathbf{e}$ ($\mathbf{e}$ contains both measurement errors and errors due to physics not contained in the model). The second sum represents a requirement for smooth time evolution, $\mathbf{m}^{\text{last}}$ being a mean of the previous mode amplitudes taken over some prior time interval T to be specified. Only a subset of L modes is used for this smoothness condition. The relative weight of the two sums is governed by the parameter α. Note that the smoothness is not applied in terms of the location of individual peaks. This would be impossible in the present application where peaks disappear and reappear with time. It is also more physical to require a smooth evolution of the gross system properties determined, instead of the peaks. The latter may fluctuate through small-scale changes in the acoustic propagation or measurement errors. Therefore, the smoothness is measured by the changes in a subset of the mode amplitudes. This subset should be the gravest modes, since they are expected to change only gradually. Higher modes describe more detailed features and may fluctuate more rapidly.

The relative weighting α^2 of the two terms in the cost function is one of the "tunable" parameters of the algorithm, as is the length T over which the mean of previous mode amplitudes is taken, and the window size. Sometimes no sensible solutions exist and even the minimized one has a large change in mode amplitudes or a large data residual. Therefore, in addition to the bound on the data residual, a maximum allowed change in mode amplitude should be set (alternatively, a maximum in the overall cost function).

TABLE 1. Example for possible identifications of three observed peaks.

	Observed peak		
	a	*b*	*c*
Possible identification	1	3	4
	0	4	5
	0	0	6

The inverse solution $\hat{\mathbf{m}}$ for the best identification is then used to make a prediction $\mathbf{G}\hat{\mathbf{m}}$ for *all* arrivals (not just the observed ones), which is superimposed on the next set of observed peaks. Windows are set again about the observed peaks, and the above procedure is repeated. The whole tracking proceeds largely automatically. Even with a poor or wrong initial condition, the method has often been observed to find the correct solution within a few time steps.

The quality (or at least consistency) of the tracking can be checked by applying the method in the forward and reverse direction, and also with different parameter settings. The rms difference between the peaks tracked in these ways can be used to tune the tracking parameters optimally. At the same time, this gives an estimate of the data uncertainty, including the (travel time) data error that results from misidentification of the peaks.

If the algorithm obviously fails and loses the peaks it is to track, it is easy to analyze why this happened. Since all identification possibilities are tried in the selection process, extra constraints or criteria in terms of simple "if" statements may be added, to prevent such occurrences. A few such extra conditions are given below for the sample application.

4. Sample application

The examples shown in this section are aimed at demonstrating the algorithm, and at the same time give typical values for the parameters used. The peak time series used in all figures comes from an acoustic tomography transmission over 52 km from a mooring pair (T3–T6) in the THETIS experiment (THETIS Group 1994). The experiment was conducted during the winter of 1991/92 in the northwestern Mediterranean convection region. Due to the winter conditions, a surface or near-surface sound channel was present. The instruments were located at 150-m depth, and all arrivals used were raylike over these distances. The arrival peaks are spread over 0.9 s in time. Only the first two or three peaks were always clearly identifiable (see

Fig. 1) and could have been tracked with pattern-following algorithms. However, some of the later arrivals also were needed to sample the shallower layers of the water column. Figure 1 and the later Fig. 3 show that there are usually more than the first three peaks and the final one visible, but that their structure and presence changes significantly. In particular, of the 11 possible arrivals selected for tracking, the later ones are too close to be resolved, that is, visible all at the same time, but subsets of those appear to be visible in different receptions.

The a priori information used in the inversion and minimization was available in the form of hydrographic profiles taken in the region, both historically and during the course of the four-month-long experiment. They allowed the estimation of the variability and vertical structure of the sound speed profiles, and a description in terms of vertical empirical orthogonal functions (EOFs). These mode amplitudes, whose expected variances were known from the analysis, form the above vector $\mathbf{m}$. From ray-tracing calculations through the mean state, 11 geometric ray arrivals were chosen for tracking, which had a spacing down to 2–3 ms in arrival time. Travel time calculations for each ray arrival with perturbed states proportional to each EOF mode yielded the observation matrix $\mathbf{G}$. For the inversion of (1), the data error was taken as 3 ms for the early arrival peaks, and 5 ms for the later ones. For more details on the analysis, see Send et al. (1995). The projection of all available hydrographic profiles onto the EOF modes showed that modes 1 and 2 seemed to have a systematic time evolution over the four months of the experiment. Therefore, the smoothness term of (2) was summed only over the first two modes, that is, $l = 1, 2$. The bounds on the modal amplitudes were taken to be 2–2.5 times the expected standard deviation of each mode. The weighting coefficient α used was of the order 0.25–0.5, and the length over which the previous mode amplitudes were averaged was up to 1.5 days (10–40 transmissions). The window size chosen was 10 or 15 ms. This resulted in several tens to thousands of possible identifications for each set of received peaks.

Figure 2 shows a typical situation in the middle of a tracking run. The top panel contains a new reception with the peaks to be identified marked with crosses. Superimposed are all the predicted locations for the arrivals to be tracked, derived from the best inverse solution of the previous reception. Solid lines indicate arrivals that were actually observed in the previous realization, and dashed lines represent the remaining predicted ones. In the example, we see that three arrivals previously not observed now seem to be present (4, 7, 8), while one has disappeared (identified as arrival 9 before). With the method explained above, the best inverse solution and corresponding peak identification is then determined. In the bottom panel, the same reception as in the top panel is redrawn, now with the best-fit solution superimposed. The predicted travel times for the identified peaks (now marked with circles) are again plotted as solid lines, while the remaining ones are shown as dashes. These will be superimposed on the subsequent reception. The last peak, which is not identified, represents some energy (and information) that is not exploited here (bottom reflected, diffracted, or sidelobe). Since no expected arrival is nearby, it is not tracked. Note that this is a way to exclude *any* unwanted peak from being tracked. The saved result of the identification of a set of peaks is not the fitted peak location (solid lines) but the observed positions of the identified peaks (circles).

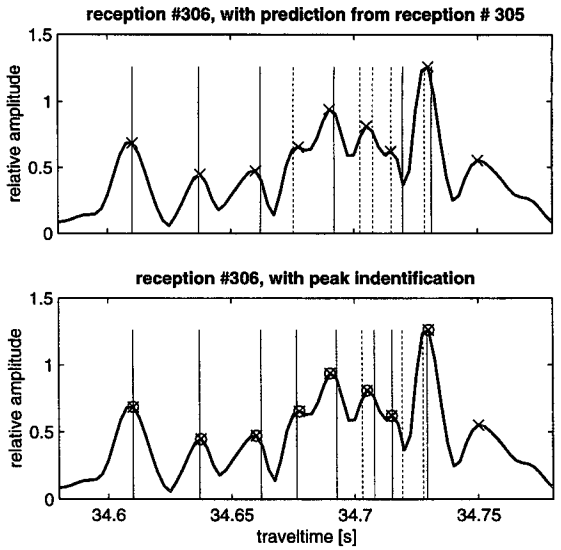

FIG. 2. Display of the identification steps in the middle of a tracking run. Top: Observed reception (heavy), with the peaks (relative maxima above a threshold) marked by crosses. Predicted possible arrivals from the previous reception are shown as vertical lines (solid: previously observed arrivals, dashed: predicted remaining possible ones). Arrivals 4, 7, and 8 were not seen before, but now seem to be present, while arrival 9 has disappeared. Bottom: Same reception as top, now with the optimum identification and the corresponding prediction (fit) indicated by vertical lines. These will be superimposed on the next reception.

Figure 3 illustrates the strongly changing nature of the received signal during two periods. It shows different possible arrivals disappear and reappear from one reception to another, and demonstrates how the algorithm can identify each observed peak with one of the expected ones. At the same time, the figure gives support to the basic assumption that all peaks can be identified—even if a new peak arises in a new position, there usually exists a theoretically possible peak with which it can be identified. Geometric ray arrivals, such as used here become very closely spaced near the final arrival. In the limit, there would then always be a predicted arrival at the location of any observed peak in this interval. While an identification in that case might

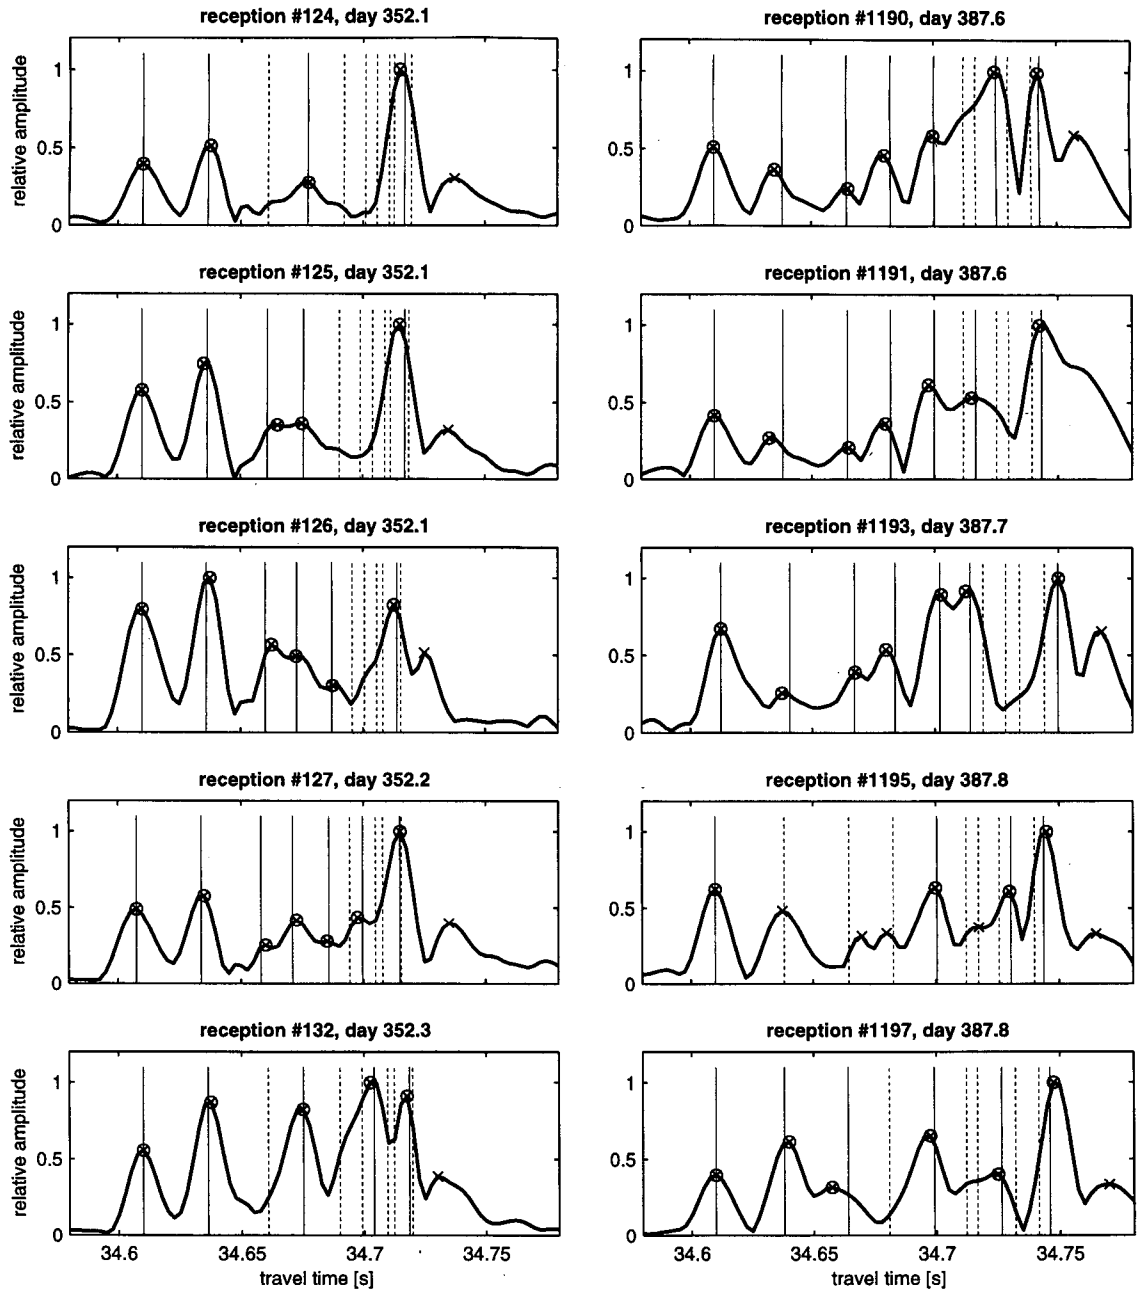

FIG. 3. Two sequences of five tomography receptions. Several peaks are seen to disappear or to be "created." Vertical lines indicate positions of the predicted arrivals from the optimum identification and inversion (solid lines are for the identified peaks). Previously not present peaks appear in locations where possible arrivals are expected. Thus, generally each observed peak could be identified, even in a strongly changing observed peak pattern.

be coincidental, such peaks would not contribute any new information to the inversion anyway.

Since the approach consists of identifying *all* observed peaks, care should be taken to avoid including too many spurious (noise) peaks, for example, by setting an appropriate height threshold for the peaks to be used. An occasional spurious peak, however, will not much affect the algorithm (either the reception is rejected since no good fit exists, or a single wrong identification is made from which the algorithm subsequently recovers).

Some extra conditions had been found necessary to restrict or enhance the set of peak identifications that are tried. Thus, it was essential that the final arrival was

among the observed peaks (however, the last peak was allowed to be omitted if it was more than 10 ms from the previous one). Due to the ordering expected in this application, the last used peak then had to be the final expected arrival. As a means of recovery from errors, the *highest* observed peak was also considered for being the main (last) arrival, since much of the time the final arrival is the highest peak.

Figure 4 shows an instant where the tracking has lost the main arrival. Correspondingly, many other peaks will also have a wrong identification. In most cases, the algorithm is able to quickly recover from such situations. When the tracking has drifted so much that the true final peak no longer contains the predicted final arrival in its window, the condition of trying the highest peak usually helps to return to the correct setting. The algorithm has occasionally shown a surprising degree of "self-correction." Figure 4 shows how the routine returns to what one would consider the correct identification, without external interfering.

The algorithm was coded in MATLAB and typical runs with 1000–2000 peak observations took a few hours on a standard workstation. Averaging forward and backward tracking runs with different parameter settings allowed identification of outliers and gave an indication of the accuracy in peak position. The result of this procedure, with subsequent calculation of daily means for each tracked peak, is shown in Fig. 5. An ordered evolution of all peaks is seen. A rough consistency check is possible from CTD stations available at intermediate times. They can be used to predict the expected peak spacing approximately. This has been superimposed on the tracking results, and generally satisfactory agreement is visible. (Complete agreement cannot be expected, especially for the late, near-surface arrivals, since the tomography data represent a 50-km average, while the CTD profiles represent point measurements in an environment that is rather variable, particularly at the surface.)

5. Conclusions

Traditional peak inversions require a three-step process: peak tracking by pure pattern-following algorithms, subsequent identification of tracked peaks, then the inversion itself. The method proposed here represents a first step toward performing all these tasks simultaneously, using the a priori information about the system. At present, several tracking runs are averaged with different parameter settings and direction. The average is then used in a final inversion.

Several improvements and future developments seem possible. For example, the restriction to geometric ray arrivals is not necessary and has in fact been relaxed in a follow-up paper (Skarsoulis et al. 1996). There, the expected arrivals (and the influence $\mathbf{G}$ of modal perturbations on them) are calculated with more advanced methods. Thus, also diffracted or modal ar-

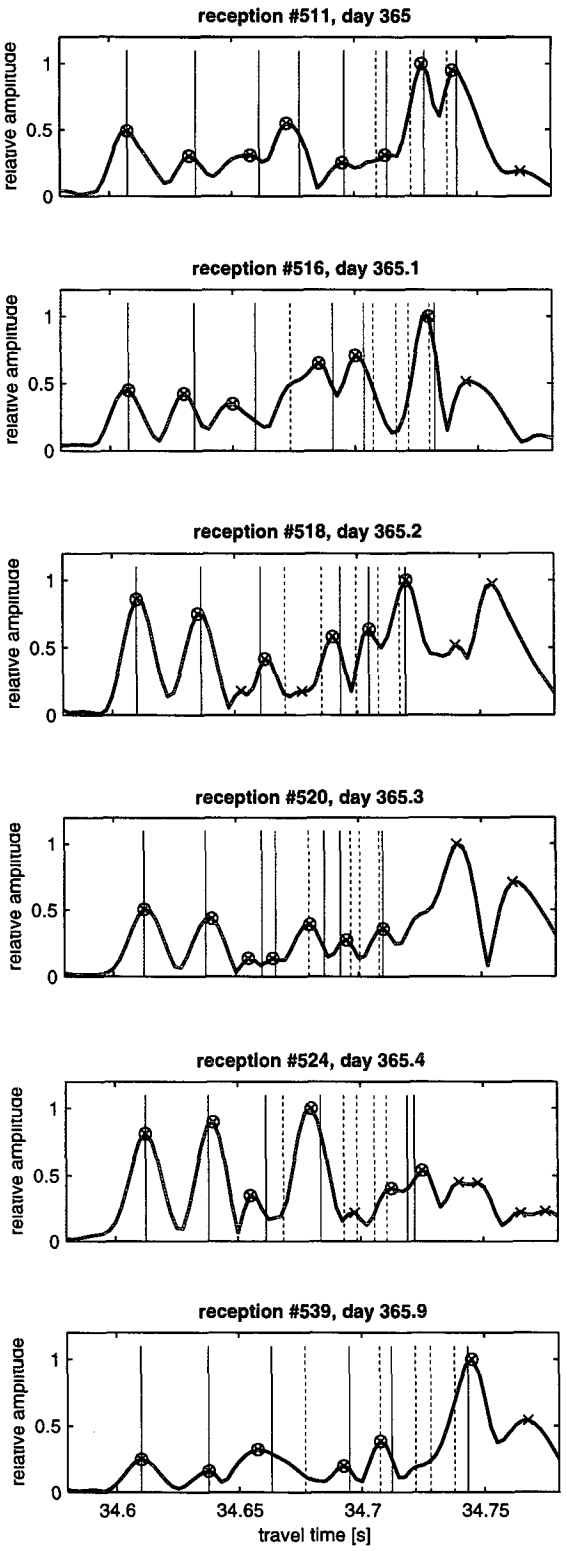

FIG. 4. Sequence of receptions in a presentation equivalent to Fig. 3. The algorithm drifts to a wrong position/identification for the main peak during a period where the actual one becomes weak or disappears. After a few time steps, the routine however returns to the likely correct identification.

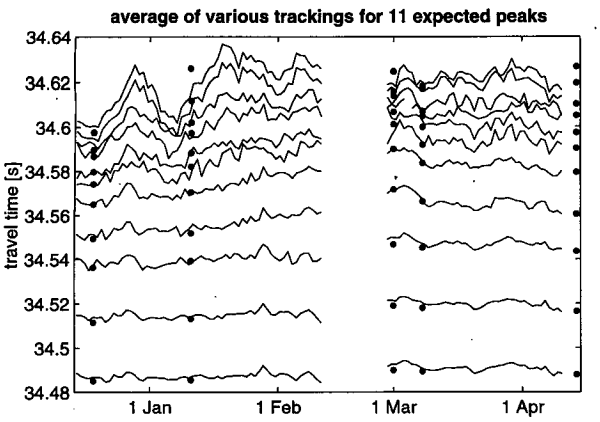

FIG. 5. Tracked peak travel times averaged over several tracking runs and over daily intervals. All expected 11 arrivals have been tracked successfully. Superimposed are predictions for the peaks from CTD data, which, however, can only be regarded as an approximate guide, since they do not fully represent horizontal averages like the tomography data.

rivals are included, which avoids the problem of expected arrivals becoming very closely spaced near the finale. New work in progress has shown also that nonlinear model-data relations can easily be accomodated in the method.

Further, one can envisage a two-pass procedure. First, the above method would be used, including the daily averaging, to obtain the overall gross evolution with some reliability. Alternatively, this could also be tried by initially tracking the first and final peaks with some other method. With that information, a more detailed pass with the proposed method is performed, where the peak positions or inverse solutions are constrained to be close to the first-pass results. Like that, high-resolution identification could be possible, without the danger of occasional drifting away or generation of outliers.

Also, the method is easily extended to include other information or constraints. For example, other than peak location, the arrival angle of peaks may be available (when using vertical arrays), for example, Howe et al. (1987). Windows would then be set around each observed peak also in angle space, and only those identifications tried which have predicted locations *and* angles within the corresponding window. In that case, the data residual in angle space would also enter the cost function.

The complete one-step identification/tracking/inversion is still difficult to implement with the above method. One problem lies in the need for some corrections *after* peak identification (mooring motion correction that depends on vertical ray launch/arrival angle). Additionally, frequently trackings from reciprocal transmissions are averaged before final inversions in order to eliminate residual clock drift. (In this case, the inversions for the *tracking* can only exploit relative peak spacing, that is, positions relative to the first peak.) These problems can be overcome. More fundamentally, the inversions for this tracking algorithm are carried out for separate instrument pairs at a time. However, many inversions are three-dimensional, using a whole instrument array at one time (e.g., Cornuelle et al. 1985). The simultaneous solution of peak identification using such 3D inversions could result in too many permutations to try (the combination of all those from single instrument pairs). This might be overcome in future by using more efficient methods for solving the minimization problem (rather than trying all possibilities), and angle information would reduce the number of possibilities significantly. The algorithm presented here, however, contains the basic principle of simultaneous minimization, even though in practice it still is a two-step process.

Acknowledgments. The THETIS experiment was a joint project by IfM Kiel, IFREMER Brest, and FORTH/IACM Crete, funded by the European Community MAST program (Contract MAST 0008-C) and the German Ministry of Science and Technology (BMFT, Contract 03F0542A). The author benefitted from interaction with B. Cornuelle during a visit to SIO, where trials with other tracking programs were made possible.

REFERENCES

Aki, K., and P. Richards, 1980: *Quantitative Seismology, Theory and Methods,* 2 Vols., W. H. Freeman and Co.

Colosi, J. A., S. M. Flatte, and C. Bracher, 1995: Internal-wave on 1000-km oceanic acoustic pulse propagation: Simulation and comparison with experiment. *J. Acoust. Soc. Amer.,* **96,** 452–468.

Cornuelle, B., and Coauthors, 1985: Tomographic maps of the ocean mesoscale. Part I: Pure acoustics. *J. Phys. Oceanogr.,* **15,** 133–152.

Howe, B. M., P. Worcester, and R. Spindel, 1987: Ocean acoustic tomography: Mesoscale velocity. *J. Geophys. Res.,* **92**(C4), 3785–3805.

Munk, W., and P. F. Worcester, 1988: Ocean acoustic tomography. *Oceanography,* **1,** 8–10.

Send, U., F. Schott, F. Gaillard, and Y. Desaubies, 1995: Observation of a deep convection regime with acoustic tomography. *J. Geophys. Res.,* **100**(C4), 6927–6941.

Skarsoulis, E. K., G. A. Athanassoulis, and U. Send, 1996: Ocean acoustic tomography based on peak arrivals. *J. Acoust. Soc. Amer.,* in press.

THETIS Group, 1994: Open-ocean deep convection explored in the Mediterranean. *EOS,* **75**(19), 217–221.