A fully nonlinear regional wave model for the Bight of Abaco
1. Nonlinear-transfer computation

Russell L. Snyder and Robert B. Long

Oceanographic Center, Nova Southeastern University, Dania, Florida

Wayne L. Neu

Department of Aerospace and Ocean Engineering, Virginia Polytechnic Institute and State University Blacksburg

Abstract. This paper is the first in a series of papers describing a fully nonlinear regional wave model for the Bight of Abaco, Bahamas. It discusses this model's hybrid representation for nonlinear transfer and the numerical errors associated with this representation. This discussion extends a number of results previously reported by *Snyder et al.* [1993], doubling both the Boltzmann integration resolution and the spectral resolution of the resulting nonlinear-transfer estimates and evaluating the errors associated with both resolutions. It also better resolves the structure of the negative midfrequency lobe of the nonlinear transfer for JONSWAP input [*Hasselmann et al.*, 1973], evaluates the errors associated with the diagnostic range of the nonlinear-transfer computation, and extends the hybrid representation to various finite depths characteristic of the Abaco Bight. It also extends the previous discussion of truncations of the hybrid representation, defining some renormalized hybrid implementations of the discrete-interaction approximation [*Hasselmann et al.*, 1985] and generalizing its selection algorithm to improve the accuracy of this approximation (at some cost in efficiency). Finally, it develops a systematic scheme for recursively selecting hybrid coefficients to define a family of recursively optimized renormalized hybrid truncations. Because this scheme selects coefficient groups (the interactions for which are scaled versions of one another or of mirror images of one another) rather than individual coefficients and because it optimizes over multiple spectral inputs, the resulting truncations perform well over a full range of peak frequency. The resulting truncations for a nominal spectral resolution of 16 prognostic and four diagnostic wave-number bands and 12 angle bands include some truncations that essentially trade a factor of 10 in efficiency for a factor of 10 in accuracy (relative to the discrete-interaction approximation). Other truncations, running only 600 times slower than the discrete-interaction approximation, give a very accurate representation of the full nonlinear transfer. These truncations, employed in sequential fashion and extended to multiple spectral resolution, enable a relatively accurate and efficient staged inverse modeling of the action-balance equation. By first inverse modeling to convergence at nominal spectral resolution, then inverse modeling to convergence at double spectral resolution (starting from the converged results for nominal spectral resolution), it should be possible to extrapolate the results of the inverse modeling with an error no worse than a few percent (assuming that the errors in the results of the inverse modeling generated by errors in the nonlinear-transfer computation are no worse than these generating errors and that the impact of other numerical errors can be similarly contained).

1. Introduction

Recent years have witnessed the appearance of a number of computer models of the action-balance equation governing the evolution of the surface-gravity-wave field. Foremost among these models is WAM, a product of many years effort on the part of its many authors. (Originally reported by the WAM Development and Implementation Group [*WAMDIG*, 1988], this model is a principal focus of the book by *Komen et al.*

[1994].) Designed primarily as an operational model to provide global wave forecasts, WAM is also variously employed as a tool for waves research. It features a coupled model of the atmospheric input that allows the evolving wave field to change the character of the wind profile driving this input [*Janssen*, 1991], a reasonable guess for the whitecap dissipation based on *Hasselmann* [1974], a highly efficient (rapidly computed) but not very accurate discrete-interaction approximation to the nonlinear transfer [*Hasselmann et al.*, 1985], and a diagnostic tail to allow for nonlinear interactions with wave components outside the prognostic range of the model. WAM also employs a semi-implicit integration scheme that provides stability for a relatively large time step, but requires a corresponding diagonalization of the functional derivative to

advance the integration (which diagonalization does not, per se, directly compromise the nonlinear-transfer computation).

It is no secret that some of the representations employed by WAM introduce significant errors into the modeling. (Published estimates of the nonlinear transfer for the discrete-interaction approximation differ significantly from estimates for EXACT-NL [*Hasselmann and Hasselmann*, 1981] (hereinafter referred to as HH). *Tolman* [1992] finds significant numerical dispersion associated with WAM's propagation algorithm.) These errors are tolerated because they are necessary to achieve the high efficiency required by an operational wave model, because they only partially distort the predictions of this model, and because the model can be tuned to partially compensate these distortions. But in the context of our inverse-modeling program, where accuracy is a prime consideration, these errors are basically unacceptable. For this reason, we are attempting to construct a model better suited to this inverse-modeling program. A critical element of the resulting model, described in this paper, is its hybrid representation for the nonlinear transfer from wave-wave interactions. A second important element, to be described in a second paper, is its discretization of the space and time variables.

In what remains of this section, we outline the inverse-modeling context for our model, develop a focus on the error of various source and propagation terms, introducing two measures of this error, and preview the remaining sections of the paper.

1.1. Context

While the Bight-of-Abaco wave model and its treatment of nonlinear transfer have general implications for other wave-modeling efforts, it is important to recognize that the development of this model has been motivated and shaped primarily by a specific ongoing long-term program to parameterize the evolutionary dynamics of the surface-gravity-wave field by comparing detailed synoptic observations of this evolution with the predictions of a fully nonlinear two-dimensional wave model [*Snyder et al.*, 1990] (hereinafter referred to as SNLdV). We next summarize this parent program.

The statistics of the wave field are described by the local action spectral density (action spectrum) $A(\mathbf{k}, \mathbf{x}, t)$, a function of vector wave number $\mathbf{k}$, horizontal position $\mathbf{x}$, and time t. The evolution of A (the spectral evolution of the wave field) is governed by the action-balance equation [*Hasselmann*, 1968; *Hasselmann et al.*, 1973], which is of the form,

$$\frac{\partial A}{\partial t} = -P + S = -P + I + N - D - B + \cdots . \qquad (1)$$

P denotes propagation terms other than $\partial A / \partial t$, and S denotes source terms, including atmospheric-input terms I, nonlinear transfer from wave-wave interactions N, dissipation from whitecapping D, and dissipation from bottom friction B.

The terms on the right-hand side of (1) determine the time rate of change of the action spectral density A. If we know how these terms depend upon A and upon other relevant observables and are given an initial state $A(\mathbf{k}, \mathbf{x}, 0)$, an appropriate set of boundary conditions, and a complete record of the observables for $t > 0$, we can in principle use this equation to calculate $A(\mathbf{k}, \mathbf{x}, t)$ for all $\mathbf{k}$, $\mathbf{x}$ in some domain, and $t > 0$.

The refraction and advection terms defining P are known functions of the gradients of A with respect to $\mathbf{k}$ and $\mathbf{x}$. The nonlinear-transfer term N is a known third-order functional of A [*Hassselmann*, 1962]. The remaining terms, however, have been only partially parameterized.

To investigate the parameterization of these terms, we are attempting to exploit the expectation that a sufficiently complete and precise record of the evolution of the action spectrum A and of those influences driving this evolution contains the information necessary to determine this parameterization. We need only (a) acquire such a record and (b) use the integral predictions of (1) to discover which combination of unknown source terms faithfully reproduces this record. To accomplish (a), we have conducted two month-long high-density high-resolution synoptic field experiments in the Bight of Abaco, a semi-enclosed section of the Bahama Banks. To accomplish (b), we will use the adjoint procedure of *Thacker and Long* [1988] to optimize the fit between model prediction and observation, adjusting the coefficients in a systematic expansion of the unknown source terms [*Snyder et al.*, 1992] (hereinafter referred to as SLL). Finally, we will explore the physical implications of the resulting best fit source terms by comparing these terms with existing theoretical predictions.

A necessary element to the success of this inverse-modeling effort is an accurate, efficient, and fully nonlinear computer model of the action-balance equation (1), specialized to the Bight of Abaco.

1.2. Focus

The success of this inverse-modeling effort will be measured primarily by the errors in the resulting estimates for the expansion coefficients and unknown source terms. These errors are in turn controlled by expansion resolution (number and type of expansion coefficients), by experiment design, and by numerical errors in the model (and adjoint-model) computations.

Ultimately, these errors must be evaluated as an important integral part of the inverse modeling of the field data . In the meantime, we are attempting to anticipate elements of this analysis in two series of papers. The first series, of which SLL is the first paper, describes some numerical inverse-modeling experiments that simulate the inverse modeling of the field data and provide a preliminary assessment of the errors associated with expansion resolution and experiment design. The second series, of which the present paper is the first paper, provides a preliminary assessment of the numerical errors of the model computation.

In discussing the numerical errors of the model computation, it is important to distinguish between known source and propagation terms and the unknown or partially known source terms to be adjusted by the inverse modeling. The form of the adjustment terms is dictated by the source expansion described by SLL; errors in the estimates for these terms are an output of the error analysis rather than an input to it. It is the errors in the numerical representation of the known source and propagation terms that contribute to these adjustment errors in the unknown source terms that are the focus of the present investigation.

These errors result from various finite choices adopted in order not to exceed the finite computer resources available to the inverse modeling. In the absence of this constraint (given unlimited computer resources), these errors and the resulting errors in the adjustment terms can in principle be made as small as desired.

Moreover, if one wants to achieve a given level of precision in an inverse-modeling computation, it is not necessary to

maintain this level of precision throughout the computation. In the early and middle stages of the computation, it is necessary only to employ a level of precision sufficient to keep the search for optimizing expansion coefficients pointed in more or less the right direction. Thus, as a practical matter, to conserve the computing resources available to the computation, it is appropriate to tune the precision implied by the current representation of known source and propagation terms to be only somewhat greater than the current precision of the unknown adjustment terms.

Our basic purpose then is to identify a hierarchy of wave models of varying precision to be employed in a sequential manner to refine the output of the inverse modeling, ultimately converging on a representation for the unknown source terms, the precision of which is as great as the available computer resources will allow, coincidentally determining the corresponding error in the resulting best fit representations for the unknown source terms.

The errors in the inversely modeled unknown source terms that result from numerical errors of the model computation are determined by propagating these numerical errors through the inverse modeling. A similar computation determines the contribution to these errors resulting from sampling errors in the field data. In this latter case, it is not possible to short-circuit this analysis because the inherent redundancy of the experimental design averages out the sampling errors, substantially complicating the relationship between these sampling errors and errors in the unknown source terms. The relationship between the numerical errors in the representations for the known source and propagation terms and errors in the unknown source terms, however, is considerably more direct. Because they are linked by (1), the signed error in these terms must sum to zero. Thus we can get some idea of the resulting errors in the unknown source terms by looking directly at numerical errors in the known source and propagation terms.

To provide a standard measure of these numerical errors, we define the proportional variances,

$$P_S \equiv \int d^2k \, (S - S_{\text{ref}})^2 \Big/ \int d^2k \, S_{\text{ref}}^2$$

and

$$P_A \equiv \int d^2k \, (A - A_{\text{ref}})^2 \Big/ \int d^2k \, A_{\text{ref}}^2, \tag{2}$$

where S_{ref} and A_{ref} are suitable reference spectra and where, depending on context, this specification focuses either on a difference of source spectra (or component thereof) or on a difference of action spectra.

Note that this specification is only one of many possible specifications (In particular, the integrands of (2) could be weighted by some power of the wave number.) and that the results of the investigation are quantitatively a weak function of this specification. Qualitatively, however, we would expect alternative specifications to result in much the same conclusions. The specification chosen is wave-number neutral in that it tends to emphasize the more energetic (active) portions of the spectrum (considered as a function of vector wave number), irrespective of where they occur.

Note also that, depending on how the reference spectra are chosen, this same specification can focus on various components of the numerical error of the nonlinear-transfer computation (in particular, on the integration, representation, and truncation errors discussed in Sections 2, 3, and 4).

We generalize the specification (2) to incorporate multiple spectral inputs by averaging this proportional variance over these inputs.

1.3. Preview

The first two papers of this series will attempt to evaluate the dependence of such measures of the error on various parameters associated with the nonlinear-transfer computation (this paper), and on various parameters associated with the discretization of the space and time variables (second paper). A third paper will address a number of other issues of concern to our model and to the success of our inverse-modeling effort.

In the next section, we begin the discussion of the nonlinear-transfer computation by briefly summarizing the hybrid scheme of *Snyder et al.* [1993] (hereinafter referred to as STHHB) and presenting some recent computations of hybrid transfer that add a diagnostic tail, double both the Boltzmann integration resolution and the spectral resolution of the previous computations, and determine the errors associated with both resolutions. This section also evaluates the errors associated with the diagnostic tail and extends the hybrid representation to a set of finite depths characteristic of the Abaco Bight.

In Section 3, we evaluate the errors associated with selected unrenormalized and renormalized truncations of hybrid transfer. These truncations include some interaction-angle truncations previously defined by STHHB, some hybrid implementations of the discrete-interaction approximation [*Hasselmann et al.*, 1985], and some hybrid implementations of the reduced-integration approximation recently proposed by *Lin et al.* [unpublished manscript]. The hybrid coefficients for these truncations are determined by various selection algorithms that focus on bin differences between interacting wave components.

In Section 4, we develop a systematic approach to the truncation of the hybrid sum. This approach selects groups of coefficients (the interactions for which are scaled versions of one another or of mirror images of one another) rather than individual coefficients and employs a calculus that makes practical the recursive optimization of the resulting truncations. It yields a multitude of renormalized truncations that are significantly more accurate than the discrete-interaction approximation at minimal cost in efficiency and defines a sequence of truncations suitable for a staged implementation of the SNLdV inverse modeling.

Finally, in Section 5, we summarize the conclusions of the investigation.

2. Errors of the Full Hybrid Representation

The nonlinear transfer interaction has been studied by a number of investigators, notably by *Hasselmann* [1962, 1963a, b, 1968] and by *Webb* [1978]. The corresponding contribution to (1) is of the form,

$$N(k, x, t) = \iiint d^2k_1 \, d^2k_2 \, d^2k_3 \tag{3}$$
$$\sigma_N \delta(k_1 + k_2 - k_3 - k)\delta(\omega_1 + \omega_2 - \omega_3 - \omega)\Pi,$$

where the spectral product Π is given by

$$\Pi(k_1, k_2, k_3, k, x, t) \equiv$$
$$A(k_1, x, t)A(k_2, x, t)\big(A(k_3, x, t) + A(k, x, t)\big)$$
$$-\big(A(k_1, x, t) + A(k_2, x, t)\big)A(k_3, x, t)A(k, x, t).$$

Here ω_i is the depth-dependent frequency corresponding to k_i, and $\sigma_N(k_1, k_2, k_3, k)$ is an interaction coefficient determined by the dynamics of the nonlinear wave-wave interaction. The delta functions express momentum and energy conservation for this interaction.

Our model employs a hybrid representation for the nonlinear transfer, previously outlined by SNLdV and described in some detail by STHHB. This representation basically incorporates a simplification first proposed by Thacker [1982] into the EXACT-NL scheme of HH, using a somewhat simpler set of integration variables. The resulting scheme retains EXACT-NL's symmetry, precision, and two-stage structure but, by transferring a spectrum-independent preintegration from the second stage to the first, dramatically accelerates the resulting second-stage computation, enabling a relatively accurate and relatively efficient representation for the nonlinear transfer. Physically, this preintegration collects together in single hybrid interactions multiple interactions belonging to identical spectral-band quadruplets. Thus all possible interactions are represented, and these interactions are represented in a uniquely efficient manner consistent with the spectral representation.

As detailed by STHHB, the hybrid representation follows directly from the expansion of A as a piecewise-constant function of k,

$$A(k, x, t) = \sum_i A_i(x, t) G_i(k) + A^R(k, x, t).$$

The basis function $G_i(k)$ is unity inside the ith spectral band and vanishes outside this band. The mean spectral density A_i is the projection of A onto G_i. A^R is a residual correction term. Similar expansions apply to P, I, N, D, and B.

Projecting (1) onto G_i yields the model equations governing the evolution of the A_i,

$$\frac{\partial A_i}{\partial t} = -P_i + I_i + N_i - D_i - B_i + \cdots. \tag{4}$$

Projecting (3) onto G_i yields the hybrid representation for N_i, presented here in its symmetrized spectral-product form,

$$N_i = \frac{1}{R_i} \sum_n T_n \sum \mu_{iq_{ns}} \Pi_{q_{ns}} + N_i^R. \tag{5}$$

R_i is the area of the ith spectral band that normalizes the projection. The outer sum is over the base interactions among quadruplet bands of the spectral representation. T_n is the positive spectrum-independent hybrid coefficient for the nth base interaction, obtained by suitably presumming various contributions to the kernel of the Boltzmann integral (3). The inner sum expands these base interactions to other angle bins of the spectral representation. The four-component vector index q_{ns} identifies the spectral bands of the expanded interaction. The coefficient,

$$\mu_{iq} \equiv \delta_{iq_4} + \delta_{iq_3} - \delta_{iq_2} - \delta_{iq_1},$$

distributes the transfer to these spectral bands. The spectral product for the expanded interaction is defined by

$$\Pi_q \equiv A_{q_1} A_{q_2} (A_{q_3} + A_{q_4}) - (A_{q_1} + A_{q_2}) A_{q_3} A_{q_4}.$$

Finally, N_i^R is a residual correction determined by A^R and the A_i.

Computation of N_i from (5) is subject to two kinds of error, integration error arising from the finite resolution of the Boltzmann integration grid (generating errors in the hybrid coefficients T_n) and representation error arising primarily from the finite resolution of the spectral grid (generating a nonzero N_i^R). The first of these errors is readily investigated by examining the convergence of the first term on the right-hand side of (5) with increasing resolution of the integration grid. STHHB evaluate this convergence for an integration resolution of 384 wave-number bands and 330 angle bands, with a linear extrapolation to infinite resolution, concluding that the resulting extrapolated integration has not yet quite converged. Comparison of the resulting hybrid estimate with that of EXACT-NL (using a piecewise-linear spectral representation), reflecting integration errors and representation errors in both computations, suggests that these estimates are comparable, with some systematic differences. Both estimates show a double peak in the negative midfrequency lobe of the nonlinear transfer for Joint North Sea Wave Project (JONSWAP) input [Hasselmann et al., 1973].

STHHB employ a spectral representation with 16 wave-number bands and 10 angle bands. The wave-number bands are all prognostic and are distributed logarithmically, with bands 1 and 16 centered on .2 Hz and 1.0 Hz, respectively. To better account for nonlinear interactions with wave components in the spectral tail beyond 1.0 Hz and to better match the angular resolution of the field measurements, the present computations employ a nominal spectral representation with 16 prognostic and four diagnostic wave-number bands and 12 angle bands. The wave-number bands are again distributed logarithmically, with bands 1 and 16 centered on .2 Hz and 1.0 Hz, respectively.

2.1. Dependence on the Resolution of the Integration Grid

As discussed by STHHB, the convergence of the transfer estimates tends toward linear as the resolution of the integration grid increases. Accordingly, it is possible to improve the results of the finite-resolution hybrid computations by linearly extrapolating these results to the case of infinite resolution. The present extrapolation employs three finite-resolution hybrid coefficient files, a higher resolution file to be extrapolated, a file with lower wave-number resolution, and a file with lower angle resolution. By decreasing the number of wave-number and angle bands in the lower resolution files by factors of 3, it is possible to guarantee that all vector indices represented in the two lower-resolution files are also represented in the higher resolution file. It follows that the dimension of the extrapolated file is the same as that of the unextrapolated higher resolution file and, more important, the extrapolated coefficients are necessarily positive.

Table 1 summarizes the four higher resolution cases included in the present convergence study for nominal spectral resolution. The indices n_k and n_ϑ are the number of wave-number and angle bands in the integration grid. The number of wave-number and angle bands in the corresponding lower resolution cases required by the extrapolations are obtained by dividing n_k and n_ϑ by 3. The index n_H is the resulting number of hybrid coefficients in the higher resolution files and in the extrapolated files. P_N is the proportional variance for an extrapolated estimate of the nonlinear transfer relative to that of the extrapolated case-4 estimate, computed from (2) for the JONSWAP spectrum,

Table 1. Integration Error for Nominal Spectral Resolution

Case	n_k	n_ϑ	n_H	$P_N^{(.25)}$	$P_N^{(.3)}$	$P_N^{(.4)}$	$P_N^{(.5)}$	$P_N^{(.7)}$
1	72	72	144,523	.2162	.1393	.1503	.0983	.1879
2	216	216	213,948	.0019	.0057	.0030	.0036	.0041
3	432	432	231,386	.0001	.0004	.0000	.0005	.0001
4	648	648	236,386	—	—	—	—	—

Here n_k and n_ϑ are the number of wave-number and angle bands in the Boltzmann integration grid, n_H is the resulting number of hybrid coefficients, and the P_N are the proportional variances between the extrapolated transfer and the corresponding extrapolated case-4 transfer for JONSWAP spectrum (6) with peak frequencies .25, .3, .4, .5, and .7 Hz.

$$A(\mathbf{k}) = \frac{1}{2}\epsilon k^{-4}\omega^{-1}e^{-\frac{5}{4}(\frac{\omega}{\Omega})^{-4}}\gamma^{e^{-\frac{1}{2\sigma^2}(\frac{\omega}{\Omega}-1)^2}}\psi(\vartheta), \quad (6)$$

with Phillips parameter $\epsilon = .01$, peak frequency Ω (nominally .3 Hz), peak enhancement factor γ (nominally 3.3), frequency spread factor $\sigma = .07$, and directional distribution,

$$\psi(\vartheta) \equiv \frac{64}{35\pi}\cos^8(\frac{\vartheta}{2}).$$

This table suggests a convergence of roughly 2% ($P_N \cong$ .0004) for case 3 and (by implication) less than 1% ($P_N \leq$ .0001) for case 4. As would be anticipated, the number of hybrid coefficients required for a fully convergent representation of the nonlinear transfer ($N_H \cong 240,000$) is somewhat larger than the corresponding number of coefficients reported by STHHB for a spectral representation with a smaller number of wave-number and angle bands. (In the course of the computations reported here, we discovered a residual redundancy in the results previously reported by STHHB. Removal of this redundancy reduces the number of coefficients for a given hybrid-coefficient file by a variable percentage that is typically of order 30% for the larger files. This reduction significantly improves the efficiency of the hybrid computation but does not alter the hybrid estimates.)

The table reveals some scatter in the proportional variances for different peak frequencies. This scatter is directly attributable to the relative phasing of the peak frequency and the frequencies of the spectral representation.

A comparison of the resulting angle-integrated JONSWAP transfer for a peak frequency of .3 Hz is given in Figure 1. This figure shows both the unextrapolated transfer (left panel) and the extrapolated transfer (right panel). In contrast to the hybrid computations presented in STHHB, the present computations employ band-averaged rather than mid-band spectral input, consistent with the interpretation of the prognostic variables A as band averages.

Comparison of Figure 1 with Figures 4 and 8 of STHHB indicates considerable agreement despite the differing spectral representations. Missing from Figure 1, however, is any direct evidence for a double peak in the negative midfrequency lobe. In fact, this lack of evidence for a double peak is somewhat coincidental. The underlying frequency width of the low-frequency lobe of this double peak is comparable with the frequency resolution of the spectral representation. Thus it is possible, with proper phasing between the peak frequency of the input spectrum and the frequencies of the spectral representation, to smooth out the appearance of this lobe. That is essentially what has happened here. The band-averaged spectral input has rendered this input less sharp (as pointed out by STHHB, midband spectral input effectively sharpens this input relative to (6)) and has slightly shifted the resulting frequency of this low-frequency lobe. These two effects completely mask the appearance of this lobe.

2.2. Dependence on the Resolution of the Spectral Grid

To better resolve the double peak in the negative midfrequency lobe for JONSWAP input and to provide a quantitative measure for the representation error of the hybrid estimates, we have extended the convergence study of Table 1 and Figure 1 to a double-resolution spectral representation with 32 prognostic and eight diagnostic wave-number bands and 24 angle bands. Table 2 summarizes the convergence of the resulting hybrid estimates with increasing resolution of the integration grid. Figure 2 shows the corresponding unextrapolated (left panel) and extrapolated (right panel) angle-integrated JONSWAP transfer.

Examination of this table and figure suggests that the convergence of the double-resolution estimates is comparable to


Figure 1. Convergence of angle-integrated hybrid transfer for nominal spectral resolution. Computed for JONSWAP spectrum (6) with peak frequency .3 Hz. Left panel shows covergence of unextrapolated hybrid estimates, right panel, convergence of extrapolated estimates: triangles, case-1 estimate; diamonds, case-2 estimate; open circles, case-3 estimate; and bullets, case-4 estimate. Units are 10^{-6} m^2 s and rad s^{-1}.

SNYDER ET AL.: NONLINEAR-TRANSFER COMPUTATION

Table 2. Integration Error for Double Spectral Resolution

Case	n_k	n_ϑ	n_H	$P_N^{(.25)}$	$P_N^{(.3)}$	$P_N^{(.4)}$	$P_N^{(.5)}$	$P_N^{(.7)}$
1	72	72	1,027,193	.1391	.0840	.1060	.0549	.0752
2	216	216	2,840,609	.0037	.0006	.0014	.0017	.0031
3	432	432	3,395,041	.0001	.0002	.0003	.0004	.0004
4	648	648	3,577,407	—	—	—	—	—

Here n_k and n_ϑ are the number of wave-number and angle bands in the Boltzmann integration grid, n_H is the resulting number of hybrid coefficients, and the P_N are the proportional variances between the extrapolated transfer and the corresponding extrapolated case-4 transfer for JONSWAP spectrum (6) with peak frequencies .25, .3, .4, .5, and .7 Hz.

that of the nominal estimates. In particular, the double-resolution estimates appear to have converged to better than 1%. The number of hybrid coefficients has increased roughly 15 times. The double peak in the negative midfrequency lobe of the angle-integrated transfer is now clearly visible, and the low-frequency lobe of this double peak is revealed to be very sharp.

Indeed the appearance of this double peak in Figure 2 suggests that the important feature may not be the double peak per se, but rather the sharpness of its low-frequency lobe. There are clearly two components to the negative midfrequency lobe, but the separation of these components into two distinct lobes may be less important than the differences in their relative characteristics. One component is broadly spread over the entire midfrequency lobe. The other component is confined to a very narrow range of frequency at the extreme low-frequency margin of this lobe.

In this context, it is important to note that the .07 frequency spread employed by both STHHB and the present computations is somewhat extreme. The mean JONSWAP spectrum has a double-sided frequency spread of .07 and .09. Thus, even though the spectral input for the present computations is less sharp than for the STHHB computations (because it is band-averaged), this input is still sharper than is implied by the mean JONSWAP spectrum. In fact, we find that with a frequency spread of .08 (not shown), the two peaks of the negative midfrequency lobe of Figure 2 become a single very sharp peak at the low frequency margin of the lobe which

(sharply) recedes at higher frequencies to what is now more of a plateau at the base of this peak than a separate peak. This suggests that it is indeed the second characterization above that is the more appropriate.

The broader issue is whether this feature, however characterized, is representative of the real evolution of the wave field. Clearly, the appearance of a double negative peak and/or very sharp low-frequency component of this negative peak for JONSWAP spectral input is related to the sharpness of this input (as determined by the peak-enhancement and frequency-spread parameters). To the extent that the JONSWAP form accurately describes the real evolution of the wave field, one might expect this evolution to reflect this feature. On the other hand, it is well known that the nonlinear transfer tends to stabilize the evolution of the spectrum. It may be no coincidence that the separation of the negative midfrequency lobe of this transfer into two distinct peaks appears to occur for an effective sharpness just slightly greater than that of the mean JONSWAP spectrum.

The hybrid estimates of Figures 1 and 2 represent the average transfer in the various spectral bands of the corresponding spectral representations. But these representations have been configured so that each band of the nominal representation contains exactly four bands of the double-resolution representation. Thus, by properly summing the double-resolution estimates, one can indirectly estimate the band averages for nominal resolution. This improved estimate of these band averages differs from a direct estimate only in that it is derived from a double-resolution representation for the input JONSWAP spectrum. Accordingly the proportional variance P_N between these two estimates provides a measure of the representation error of the nominal estimate.

Table 3 shows this proportional variance as a function of peak frequency, and Figure 3 compares these two estimates and the EXACT-NL estimate from Figure 8 of STHHB for a peak frequency of .3 Hz.

As before, there is substantial agreement among the three estimates, but there are systematic differences. Furthermore, the differences remain pretty much as reported by STHHB. The biggest difference is that the EXACT-NL estimate does not descend quite as sharply from its positive low-frequency lobe to its negative midfrequency lobe as do the hybrid es-


Figure 2. Convergence of angle-integrated hybrid transfer for double spectral resolution. Computed for JONSWAP spectrum (6) with peak frequency .3 Hz. Left panel shows convergence of unextrapolated estimates, right panel, convergence of extrapolated estimates: diamonds, case-2 estimate; open circles, case-3 estimate; and bullets, case-4 estimate. Units are 10^{-6} m^2 s and rad s^{-1}.

Table 3. Representation Error for Nominal Spectral Resolution

Ω	P_N
.25	.0129
.3	.0113
.4	.0114
.5	.0119
.7	.0127

Here Ω is the peak frequency in Hz, and P_N is the corresponding proportional variance between the extrapolated nominal-resolution case-4 estimate and the averaged extrapolated double-resolution case-4 estimate, computed for JONSWAP spectrum (6).

timates. Because the descent of the EXACT-NL estimate is quite similar to that of the unextrapolated lower-integration-resolution case-1 and case-2 hybrid estimates of Figure 1, we believe this difference is primarily the result of the incomplete convergence of the EXACT-NL estimate. A second qualitative difference is that the maximum in the positive high-frequency lobe of the EXACT-NL estimate occurs at higher frequency. The improved hybrid estimate lowers the amplitude of the high-frequency lobe of the nonlinear transfer relative to that of the direct estimate, but it does not reproduce the dimple of the EXACT-NL estimate or appreciably shift the frequency of the maximum. We do not presently understand this difference.

Table 3 suggests that for nominal spectral resolution the representation error of the extrapolated case-4 hybrid estimates is consistently of order 11% over an entire range of peak frequency. This error is considerably larger than the integration error, suggesting that nominal spectral resolution may be insufficient for the final stages of the SNLdV inverse-modeling computation and raising the question, what spectral resolution is sufficient for these stages? In order to provide at least a partial answer to this question, we have pushed the computation still one step further.

Table 4 shows the proportional variance for the JONSWAP input (6) with peak frequency .3 Hz for eight choices of spectral resolution. In each case, we have compared the unextrap-


Figure 3. Comparison of various estimates for the angle-integrated transfer for JONSWAP spectrum (6) with peak frequency .3 Hz: diamonds, EXACT-NL estimate; open circles, extrapolated nominal-spectral-resolution case-4 estimate; and bullets, averaged extrapolated double-spectral-resolution case-4 estimate. Units are 10^{-6} m^2 s and rad s^{-1}.

olated case-2 computation with an unextrapolated quadruple-spectral-resolution case-2 computation with 64 prognostic and 16 diagnostic wave-number bands and 48 angle bands. All computations are collapsed to a nominal spectral resolution of 16 prognostic and four diagnostic wave-number bands and 12 angle bands before evaluating the proportional variance. (To simplify this computation, we have employed case-2 rather than case-4 integration resolution. We have not extrapolated the results of this computation. We believe that an extrapolated case-4 computation would have produced similar results.)

This table increases the estimate for the representation error at nominal spectral resolution from 11% to 13% and suggests that the representation error of a double-spectral-resolution computation is of order 4% while the representation error of a computation with quadruple wave-number resolution and double angle resolution is of order 2%. Clearly, it is not possible to match the 1% case-4 integration error without increasing the spectral resolution roughly fourfold (or, as discussed below, in some way extrapolating these estimates).

Figure 4 shows the corresponding angle-integrated transfer for nominal, double, and quadruple spectral resolution (main diagonal of Table 4), collapsed to a nominal spectral resolution of 20 total wave-number bands and 12 angle bands (left panel) and to a double spectral resolution of 40 total wave-number bands and 24 angle bands (right panel).

The visual contrast between the left-hand and right-hand panels of Figure 4 might lead one to conclude that nominal spectral resolution is not sufficient for wave modeling in general or for the inverse-modeling program of SNLdV in particular. We would emphasize, however, that, regardless of the spectral resolution, the prognostic equations (4) still govern the corresponding band averages. What is important to the program of SNLdV is not that these averages resolve all variations in nonlinear transfer with frequency and angle, but that the associated representation error be small enough to ensure small errors in the estimates for the unknown source terms. Visually, what is important in this figure is not the difference between the nonlinear transfer depicted in the left-hand and right-hand panels, but rather the differences between the nominal-, double-, and quadruple-spectral-resolution estimates of the left-hand panel. Ultimately, the question of what spectral resolution is adequate for the inverse-modeling program of SNLdV must be answered by explicitly evaluating the resulting errors in the inversely modeled source terms. Such an evaluation is beyond the scope of the present paper.

Table 4. Representation Error for Extended Spectral Resolution

T	$P_N^{(12)}$	$P_N^{(24)}$	$P_N^{(48)}$
16	.0162	.0078	.0061
32	.0080	.0015	.0008
64	.0068	.0004	—

Here the P_N are the proportional variances between various resolution unextrapolated case-2 estimates of N and the corresponding unextrapolated case-2 estimate for quadruple spectral resolution. Computed for JONSWAP spectrum (6) with peak frequency .3 Hz for 16, 32, and 64 prognostic wave-number bands T and 12, 24, and 48 angle bands S. Spectra collapsed to nominal spectral resolution before computing proportional variance. The central proportional variance provides an estimate for the representation error for double spectral resolution.

Figure 4. Comparison of unextrapolated case-2 angle-integrated hybrid transfer for various-resolution JONSWAP spectra (6) with peak frequency .3 Hz. Displayed transfer collapsed to nominal spectral resolution (left panel) and to double spectral resolution (right panel): diamonds, nominal-spectral-resolution estimate; open circles, double-spectral-resolution estimate; and bullets, quadruple-spectral-resolution estimate. Units are 10^{-6} m^2 s and rad s^{-1}.

However complex the relationship between the numerical errors of the nonlinear-transfer computation and the resulting errors in the inversely modeled source terms, it is important to note that these resulting errors can be significantly reduced by first inverse modeling at nominal spectral resolution, then inverse modeling at double spectral resolution and extrapolating the results of the inverse modeling to the case of infinite spectral resolution (in much the same way that the the error of the Boltzmann integral is reduced by extrapolating the Riemann approximant to this integral to the case of infinite integration resolution). To the extent that the resulting errors are comparable with the numerical errors of the nonlinear transfer estimates, one would expect the errors of the extrapolated results to be no greater than a few percent.

Another way of decreasing the representation error of the hybrid computation is to extend this computation to a piecewise-linear spectral representation. As pointed out in appendix A of STHHB, however, a consistent application of the piecewise-linear representation requires a deconvolution of the resulting nonlinear-transfer estimates. In the absence of this deconvolution, it is not clear that the representation error of the resulting estimates or of spectral evolutions employing these estimates is significantly improved relative to the piecewise-constant computation.

For this reason and because of the inherent simplicity of the piecewise-constant computation (and inherent complexity of a deconvolved piecewise-linear computation), increasing the spectral resolution of the piecewise-constant computation may provide a more practical approach to improving the representation error than extending this computation to a piecewise-linear spectral representation at the same resolution. Ultimately, the relative error of these two approaches needs to be directly evaluated by implementing a piecewise-linear extension to the hybrid computation.

2.3. Dependence on Diagnostic Wave Components

The nonlinear character of the nonlinear transfer interaction implies that interactions involving wave components external to the prognostic range of the wave-model computation can contribute significantly to the nonlinear transfer within this range. In particular, interactions involving wave compo-

nents with frequencies somewhat larger than the upper limit frequency of the wave model can contribute significantly to this transfer. To allow for these interactions, we have followed the WAM prescription of attaching a number of diagnostic wave-number bands to the model (nominally four such bands for nominal spectral resolution), extrapolating A in these bands from the uppermost prognostic wave-number band (assuming a Phillips tail) and incorporating these bands into the nonlinear-transfer computation.

Similarly, interactions involving wave components with frequencies somewhat below the lower limit frequency of the wave model can also contribute to the modeled transfer.

In order to evaluate the dependence of the hybrid transfer on such interactions, we have extended our nominal-spectral-resolution computations (with extrapolation) to include various numbers of wave-number bands above and below the prognostic range of the wave model. (Note that the important variable is not the number of bands per se but the frequency range represented by this number of bands.) To simplify these computations, we have employed case-2 rather than case-4 integration resolution. Table 5 shows the resulting dependence of the proportional variance P_N, computed from (2) over the prognostic range of the spectrum, between the nonlinear transfer for various numbers of upper and lower external bands and the nonlinear transfers for a large number (24) of upper external bands or a large number (10) of lower external bands. These proportional variances are computed for the JONSWAP spectrum (6) with peak frequency .25, .3, .4, .5, and .7 Hz.

The upper section of Table 5 suggests that, at the lower peak frequencies, four upper external diagnostic bands are sufficient to represent the nonlinear transfer with 1% accuracy. At higher peak frequencies, however, the error increases to 3-4% at .5 Hz and 13% at .7 Hz. Doubling the number of diagnostic bands to eight extends the range over which the error is less than or equal to 1% to a peak frequency of .5 Hz. While a modest increase in the number of diagnostic bands does not severely burden the computation, particularly if implemented adaptively, the early development of the spectrum is nonetheless somewhat problematic for the nonlinear transfer computation (as it is in a number of other important respects).

Table 5. Errors Resulting From the Finite Range of the Spectral Representation

L	U	n_H	$P_N^{(.25)}$	$P_N^{(.3)}$	$P_N^{(.4)}$	$P_N^{(.5)}$	$P_N^{(.7)}$
0	0	137,722	.0005	.0011	.0118	.0689	.6153
0	4	213,948	.0001	.0001	.0004	.0012	.0177
0	8	279,263	.0000	.0000	.0001	.0001	.0007
0	14	336,979	.0000	.0000	.0000	.0000	.0000
0	24	371,504	—	—	—	—	—
0	4	218,421	.0006	.0000	.0000	.0000	.0000
1	4	239,529	.0000	.0000	.0000	.0000	.0000
10	4	397,149	—	—	—	—	—

Here L is the number of lower external wave-number bands, U is the number of upper external wave-number bands (number of diagnostic bands), n_H is the resulting number of hybrid coefficients, and the P_N are the proportional variances between the extrapolated nominal-spectral-resolution case-2 transfer and the reference transfer defined by the final entry in each section of the table, computed for the JONSWAP spectrum (6) with peak frequencies .25, .3, .4, .5, and .7 Hz.

The lower section of Table 5 suggests that so long as the lower limit frequency of the model is somewhat lower than the lowest anticipated peak frequency, there is no need for additional lower external bands. If such a need develops, it is probably most readily accommodated, not by attaching lower external bands, but by simply extending the prognostic range of the model.

2.4. Dependence on Depth

The hybrid computations reported above assume deep-water dispersion. These computations need to be amended to allow for the finite depth of the Abaco Bight.

At any location in the Bight the spectral range divides into two parts, a low-wave-number (low-frequency) part that feels the bottom and a high-wave-number (high-frequency) part that does not. This division depends on the wave number k relative to a threshold wave number k_H determined by the local depth H. For $k \leq k_H$, the dispersion is depth-dependent; for $k > k_H$, this dispersion is not.

A similar division applies to the quadruplet interactions that contribute to the hybrid sum. If all of the interacting wave numbers for a given interaction are larger than k_H, then the appropriate hybrid coefficient is a deep-water coefficient, estimated from a Boltzmann integration with deep-water dispersion. If, on the other hand, any of these interacting wave numbers are smaller than or equal to k_H, this coefficient must be estimated from a Boltzmann integration with depth-dependent dispersion.

Accordingly, we may split the nonlinear-transfer computation into two parts. One part is defined by a single appropriately ordered deep-water hybrid coefficient file and is summed from a start index depending on depth to an end index depending on the number of diagnostic bands. The other part is defined by one of several finite-depth hybrid coefficient files and is summed to an end index depending on the number of diagnostic bands. We simplify the computation of the depth-dependent part of this transfer by ordering the interior grid points of the model by increasing depth. This allows neighboring points to share the same hybrid coefficient files.

In fact, as will be discussed in more detail in the second paper in this series, this depth splitting is complicated by a second time splitting, introduced to reduce the burden of the nonlinear-transfer computation within the wave model. This time splitting is related to an instability, associated with the nonlinear transfer, that first appears in the uppermost prognostic wave-number band. To avoid this instability, it is necessary to employ a relatively small time step. It is not necessary, however, to advance all nonlinear interactions at this same time step. Interactions among low-wave-number components can be advanced at multiples of the time step without causing the computation to become unstable. To take advantage of this result, we further subdivide the hybrid coefficient files and embed the main time-stepping loop of the computation in one or more outer loops in which we separately evaluate the nonlinear transfer from low-wave-number interactions.

The computation of depth-dependent hybrid coefficients is more consuming of computer resources than is the computation of deep-water coefficients. In particular, the evaluation of the interaction coefficient σ_N in the integrand of (3) is more complex, and the integration cannot take advantage of $\hat{k}$ scaling (as employed by Hasselmann [1963a, b], Resio and Perrie [1991], and STHHB). Nonetheless, the dependence of this interaction coefficient on depth is known [Hasselmann, 1962; Herterich and Hasselmann, 1980], and, because of the depth splitting, it is possible to bypass considerable computation by first checking whether a given interaction is in fact depth-dependent before evaluating its interaction coefficient (retaining only interactions that are depth-dependent). We have in fact succeeded in extending the deep-water computation to a set of five depths characteristic of the Abaco Bight, 2, 4, 6, 8, and 10 m, computing extrapolated case-4 hybrid coefficient files for all five depths. The corresponding nominal-spectral-resolution threshold-wave-number bins for these computations are bin 13, 10, 8, 7, and 6. The integration and representation errors for these finite-depth representations are presumed comparable with the corresponding errors for the deep-water representation.

2.5. Triplet Interactions

It has been suggested that in shallow water it is possible for resonant nonlinear interactions among triplets of interacting wave components to contribute to the evolution of the wave field [Beji and Battjes, 1993]. Our treatment of nonlinear transfer does not allow for such interactions. Could this neglect significantly impact the SNLdV inverse modeling? We think not. The Abaco Bight is shallow (relative to the wavelengths present) only along its periphery. Whatever the wind direction, this periphery divides into two parts, an upwind periphery and a downwind periphery. Along the upwind periphery, kH is restricted to larger values by the small amplitude of longer wave components at short fetch. Here there can be no significant triplet interactions. Along the downwind periphery, the products of the nonlinear interactions, whether quadruplet or triplet, are primarily absorbed at the boundary without crossing into deeper water. Because the observation stations are all located in deeper water, these products cannot significantly impact the inverse modeling.

3. Selected Truncation of the Hybrid Representation

The rather severe computational demands of the nonlinear-transfer computation, exacerbated by the necessity to estimate this transfer at each spatial grid point and each time

step of a wave-model integration, has prompted a number of approximate schemes for calculating this transfer. Typically, these schemes involve either a less-than-convergent Boltzmann integration, explicit truncation of the resulting Riemann sum, renormalization of the resulting transfer, or some combination of the three.

It should be noted that, in most schemes for calculating nonlinear transfer, the accuracy of the scheme is inversely related to its efficiency. The more efficent the scheme, the less accurate it can be and vice versa. The hybrid scheme is in some respects an exception to this rule. In other schemes, the number of terms in the Riemann sum increases without limit as the integration resolution and accuracy increase. In the hybrid scheme, this number approaches an asymptotic limit. Because the hybrid coefficients are computed only once, there is little reason not to employ the most convergent set of coefficients that can practically be estimated, currently the extrapolated case-4 set of Section 2. While the resulting full hybrid sum may be too long for some applications, it provides a excellent point of departure for subsequent truncation and renormalization. In effect, the hybrid scheme eliminates integration error from the resulting estimate of nonlinear transfer, leaving only the representation error and the error associated with subsequent truncation and renormalization. Ultimately, these errors are determined by the resolution of the spectral representation and by the number and type of hybrid coefficents retained.

STHHB discuss a number of unrenormalized truncations that are effective in reducing the dimension of the hybrid sum without significantly degrading the resulting estimate for the nonlinear transfer. Particularly effective among these are the relatively input-nonspecific interaction-angle truncations that eliminate coefficients for which the angle bins for any two of the four interacting wave components differ by more than some threshold value and the input-specific size-ordered truncations that eliminate coefficients that contribute only marginally to the nonlinear transfer specific to particular input spectra. A particular combination of these two truncations, retaining only 3 to 6% of the coefficients (▷ and ◇ of Figure 12 of STHHB), gives a good to excellent representation for the nonlinear transfer associated with the JONSWAP input (6), albeit for a single specific choice of peak frequency.

Probably the most efficient approximation to the nonlinear transfer is the discrete-interaction approximation [*Hasselmann et al.*, 1985]. Employed to represent the nonlinear transfer in WAM, this approximation allows only two nonlinear interactions per spectral band, an efficiency roughly equivalent to that of a hybrid sum with $2T$ terms, where T is the number of wave-number bands in the spectral representation. (An additional factor of S, the number of angle bands in the spectral representation, is expressed as an inner loop in the hybrid sum (5).) The strength of the discrete interactions, however, must be dramatically increased to give a (marginally) acceptable estimate for the resulting nonlinear transfer. (It is as if the interactions that are not explicitly included in the discrete-interaction approximation, and that together define the bulk of the transfer, sum together to give a contribution to this transfer much larger than but with a shape similar to that of the allowed interactions. What is remarkable is that so crude an approximation can be tuned to give a reasonably good account of the evolution of several critical spectral parameters (total energy and peak frequency).

Let k_1, k_2, k_3, and k_4 be the vector wave numbers of the wave components for a given nonlinear interaction, and let $\vartheta_i \equiv \arg(k_i)$ be the corresponding directions of propagation for these wave components. Symmetrization of the Boltzmann integral [HH] restricts the range of these variables to $k_1 \leq k_2$, $k_3 \leq k_4$, $\vartheta_1 \leq \vartheta_3 \leq 2\pi - \vartheta_1$, and $\vartheta_2 \leq \vartheta_4 \leq 2\pi - \vartheta_2$. In this context, the discrete-interaction approximation in effect truncates the nonlinear transfer by restricting the integration space of a three-dimensional version of the HH-symmetrized Boltzmann integration to the immediate neighborhood of a single point (and its mirror image) for which $k_1 = k_2$ and for which ϑ_3 and ϑ_1 differ by 33.6 degrees.

Recently, R. Q. Lin et al. (unpublished manuscript, 1996) have proposed a reduced-integration approximation to the Boltzmann integral that restricts this integration space to a broad neighborhood of the line defined by $k_3 = k_1$ (and $k_4 = k_2$), along which the coupling coefficient $\sigma_N(k_1, k_2, k_3, k_4)$ has an integrable singularity. (HH respond to this same singularity by stretching the EXACT-NL integration variables to provide a greater concentration of integration points in the neighborhood of this line.) Note, however, that the spectral product vanishes identically for $k_3 = k_1$. It follows that the immediate neighborhood of this line does not contribute substantially to the integral. Therefore the reduced-integration approximation requires a significantly broader neighborhood than might otherwise be required.

Because the vector indices for the hybrid coefficients directly identify the spectral bins of the four interacting wave components, it is relatively easy to define hybrid implementations for both the discrete-interaction approximation and the reduced-integration approximation. One simply takes a full set of hybrid coefficients and eliminates those with vector indices that do not in some approximate sense satisfy the conditions for the approximations.

Note that the resulting hybrid implementations of the discrete-interaction approximation and reduced-integration approximation are only roughly equivalent to the original approximations. In each case the hybrid implementation in effect adds to the interactions of the approximation a neighboring range of interactions defined by the spectral resolution. Because of the presumming, it is not possible to eliminate the neighboring interactions. However, the spirit of the approximation is retained, and the additional interactions should not adversely affect either the accuracy of the truncation or its efficiency.

We next evaluate the truncation error of a series of selected hybrid truncations, including six renormalized implementations of the discrete-interaction approximation and generalizations of this approximation DI1 through DI6, two implementations of the reduced-integration approximation RI1 and RI2, with and without renormalization, and two interaction-angle truncations AT2 and AT3, with and without renormalization.

The discrete-interaction implementations attempt to match the wave-number bins and angle bins of the hybrid coefficients to the specifications for the discrete-interaction approximation. All of these implementations require the angle bins for ϑ_1, ϑ_2, and ϑ_4 to be identical and require the angle bin for ϑ_3 to differ by one. The low-level discrete-interaction implementations DI1 through DI3 require the wave-number bins for k_1 and k_2 to be identical. DI4 allows the bins for k_1 and k_2 to differ by no more than 2, DI5, by no more than 4, and DI6, by no more than 6. DI1, DI2, and DI3 all require the bin for k_4 to be 2 larger than the bin for k_1. DI1 requires the bin for k_3 to be 2 smaller than the bin for k_1. DI2 requires the bin for k_3 to be no more than 2 smaller than the bin for k_1, DI3,

no more than 3 smaller. DI4, DI5, and DI6 require the bin for k_4 to be no more than 3, 5, and 7 larger than the bin for k_1 and require the bin for k_3 to be no more than 3, 5, and 7 smaller than the bin for k_1.

The reduced-integration implementation RI1 restricts the integration space for the approximation as much as it can be restricted within the context of the hybrid scheme. This implementation requires the angle bins for ϑ_1 and ϑ_3 to differ by no more than 1 and the wave-number bins for k_3 and k_1 to differ by no more than 1. The reduced-integration implementation RI2, a less restrictive implementation more compatible with the integration limits of the original approximation, requires these wave-number and angle bins to differ by no more than 4 and 2, respectively. Neither implementation allows these differences to simultaneously vanish (because the corresponding spectral product is identically zero).

The AT2 and AT3 interaction-angle truncations require the angle bins for ϑ_1, ϑ_2, ϑ_3, and ϑ_4 to differ by no more than 2 and 3, respectively.

Note that these discrete-interaction, reduced-integration, and interaction-angle truncations are resolution-dependent. A given truncation for nominal spectral resolution incorporates a somewhat different set of fundamental interactions (and gives a somewhat different estimate for the nonlinear transfer) than does the same truncation for double spectral resolution. The results presented in this section and in Section 4 are for nominal spectral resolution.

3.1. Renormalization

The renormalization of a given truncation or approximation can be accomplished in several ways. WAM essentially tunes the discrete-interaction approximation to match the corresponding evolution of total energy and peak frequency to that predicted by EXACT-NL. We adopt two somewhat simpler and more direct procedures for determining an optimum renormalization, related to the differential and integral measures of the error defined by (2). In either case, we employ reference spectra defined by an extrapolated case-4 computation of nonlinear transfer rather than by an EXACT-NL computation. Adopting the JONSWAP input spectrum (6) and minimizing the corresponding proportional variance,

$$\int d^2k \, (\mu_N N - N_{\rm ref})^2 \Big/ \int d^2k \, N_{\rm ref}^2,$$

gives the transfer-based minimum proportional variance P_N and minimizing renormalization factor μ_N,

$$P_N = 1 - \left(\int d^2k \, N N_{\rm ref}\right)^2 \Big/ \left(\int d^2k \, N^2 \int d^2k \, N_{\rm ref}^2\right)$$

and

$$\mu_N = \int d^2k \, N N_{\rm ref} \Big/ \int d^2k \, N^2.$$

This result is readily generalized to incorporate multiple input spectra (in which case P_N is the minimum mean proportional variance). In either case, P_N and μ_N are functions of the JONSWAP parameters defining the input spectra.

Similarly, adopting a standard specification for the atmospheric input and case-2 dissipation, employing a stable and convergent time step, multiplying the hybrid coefficients by the renormalization factor μ_A, integrating the homogeneous prognostic equations from a null spectrum, and minimizing the resulting proportional variance at some prescribed stage of the evolution (or resulting mean proportional variance at mul-

tiple prescribed stages of the evolution) gives the evolution-based minimum proportional variance P_A and minimizing renormalization factor μ_A. P_A and μ_A are (weak) functions of the parameters defining the atmospheric input and dissipation.

The integration of the model equations will be discussed at length in the second paper of this series. Suffice to say, the present integration of homogeneous deep-water equations to estimate P_A for a given truncation of hybrid transfer employs an explicit forward-differencing scheme and, to maintain stability, a relatively small time step (12 s). Atmospheric input and dissipation are defined by case 4 of SLL, with $M_\alpha = M_\beta = M_\gamma = 1$ and $W = 10$ m s^{-1}. The corresponding dimensionless coefficients Q_1 (10^{-6}), Q_2 (.2), and Q_3 (.004), expressing the strength of the Phillips [1957] component of the atmospheric input, Miles [1957] component of the atmospheric input, and a second-order version of Hasselmann [1974] dissipation, respectively, are chosen to give a reasonable four-hour evolution of the spectrum. (Note that Q_1 is relative to a Phillips interaction that goes as the second power of the air/water density ratio rather than as the first power of this ratio originally incorporated into the SLL relations.)

A detailed comparison of the resulting evolution with a comparable real evolution is not important to the discussion. What is important is the extent to which a given truncation of the hybrid transfer reproduces the evolution for full hybrid transfer.

3.2. Results

Table 6 shows the minimizing renormalization factors μ_N and μ_A and corresponding minimum proportional variances P_N and P_A for multiple-input optimizations of an enlarged set of truncations of hybrid transfer, including discrete-interaction truncations DI1 through DI6, reduced-integration truncations RI1 and RI2, and interaction-angle truncations AT2

Table 6. Renormalization of Selected Truncations

Case	n_H	μ_N	P_N	μ_A	P_A
DI1	32	47.8	.3655	48.7	.3573
DI2	64	35.3	.3100	67.5	.3653
DI3	94	36.4	.1560	26.2	.2106
DI4	650	6.20	.1187	6.84	.0881
DI5	1,126	4.50	.0605	4.90	.0398
DI6	1,376	3.98	.0566	4.30	.0451
RI1	34,583	—	.1987	—	.1651
RI1	34,583	1.38	.1337	1.84	.0600
RI2	91,650	—	.0084	—	.0733
RI2	91,650	1.03	.0077	0.78	.0539
AT2	33,051	—	.0034	—	.0023
AT2	33,051	1.03	.0024	1.05	.0018
AT3	65,493	—	.0002	—	.0002
AT3	65,493	1.00	.0002	1.00	.0002

Here μ_N and P_N are the optimum transfer-based renormalization and mean proportional variance for JONSWAP spectrum (6) with peak frequencies .3, .4, .5, .6, and .7 Hz, μ_A and P_A are the optimum evolution-based renormalization and mean proportional variance for a 4-hour homogeneous evolution of the spectrum, evaluated at hourly intervals, and n_H is the number of hybrid coefficients in the truncation. DI truncations are hybrid implementations of the discrete-interaction approximation, RI truncations, hybrid implementations of the reduced-integration approximation, and AT truncations, interaction-angle truncations. Several truncations are unrenormalized.

and AT3. The transfer-based optimization employs multiple JONSWAP input spectra (6) with peak enhancement 3.3 and peak frequencies .3, .4, .5, .6, and .7 Hz. The evolution-based optimization compares the action spectra for hours 1 through 4 of a homogeneous evolution, starting from the null spectrum. Note the similarity of results for the two optimizations. Except in the case of the RI2 truncation, the renormalization factors are generally comparable, and so are the corresponding proportional variances. Note also the relatively high level of proportional variance for the discrete-interaction and reduced-integration truncations.

Figures 5, 6, and 7 show the corresponding angle-integrated transfer, differential transfer, and homogeneous evolution of the angle-integrated spectrum for full hybrid transfer and for the AT2, DI5, and DI1 renormalized truncations of this transfer.

It is clear from these figures that both the AT2 and DI5 truncations give a considerably more realistic representation for the nonlinear transfer and for the homogeneous evolution of the spectrum than does the discrete-interaction approximation (as represented by the DI1 truncation), without too dramatically increasing the computation time. Of particular interest are the evolutions for the DI5 and DI1 truncations shown in Figure 7. The DI1 evolution exhibits the pre-

mature low-frequency spreading and coincident flattening of the angle-integrated spectrum characteristic of the discrete-interaction approximation. The DI5 evolution is considerably more realistic. The principal failing of this evolution is that the level of the high-frequency tail is somewhat larger than that of the Phillips equilibrium. Coincidentally, as evidenced by Figure 6, the angular dependence of the transfer in this tail is qualitatively wrong.

Because it involves the same number of fundamental interactions (which are, however, interactions among spectral bands), the efficiency of the DI1 truncation is necessarily comparable with that of the discrete-interaction approximation. Therefore, the DI5 truncation can be expected to be roughly 35 times slower than this approximation.

For an additional one to two order-of-magnitude increase in computation time, it is clear that the interaction-angle truncations AT2 and AT3 result in a truncation error comparable with (smaller than) the computational errors of Section 2. While these truncations must be handled with care (because they depend to some extent on the angular localization of the spectrum, as do all the truncations of Table 6), they provide a very accurate estimate of this transfer well short of the full hybrid sum. Additionally, these truncations require only modest renormalizations of order 1.
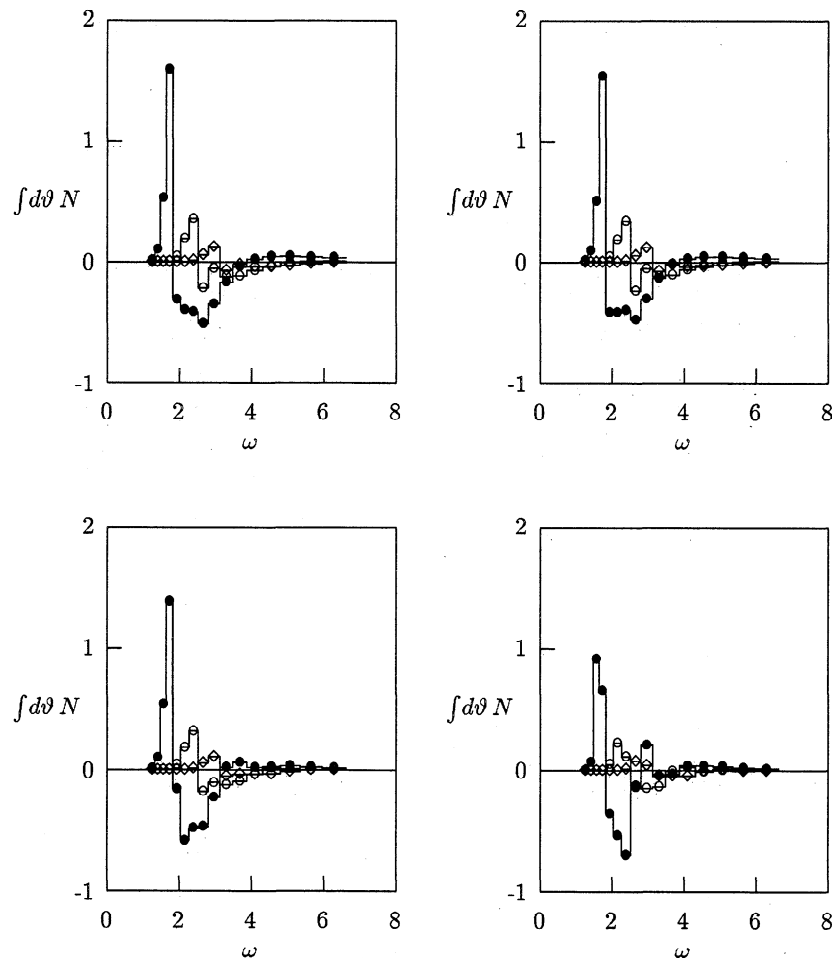

Figure 5. Angle-integrated transfer for various renormalized discrete-interaction truncations of Table 6. Computed for JONSWAP spectrum (6) with peak frequency .3 Hz (bullets), .4 Hz (open circles), and .5 Hz (diamonds). Upper left panel is full hybrid transfer. Upper right panel is AT2 truncation. Lower left panel is DI5 truncation. Lower right panel is DI1 truncation. Units are 10^{-6} m^2 s and rad s^{-1}.
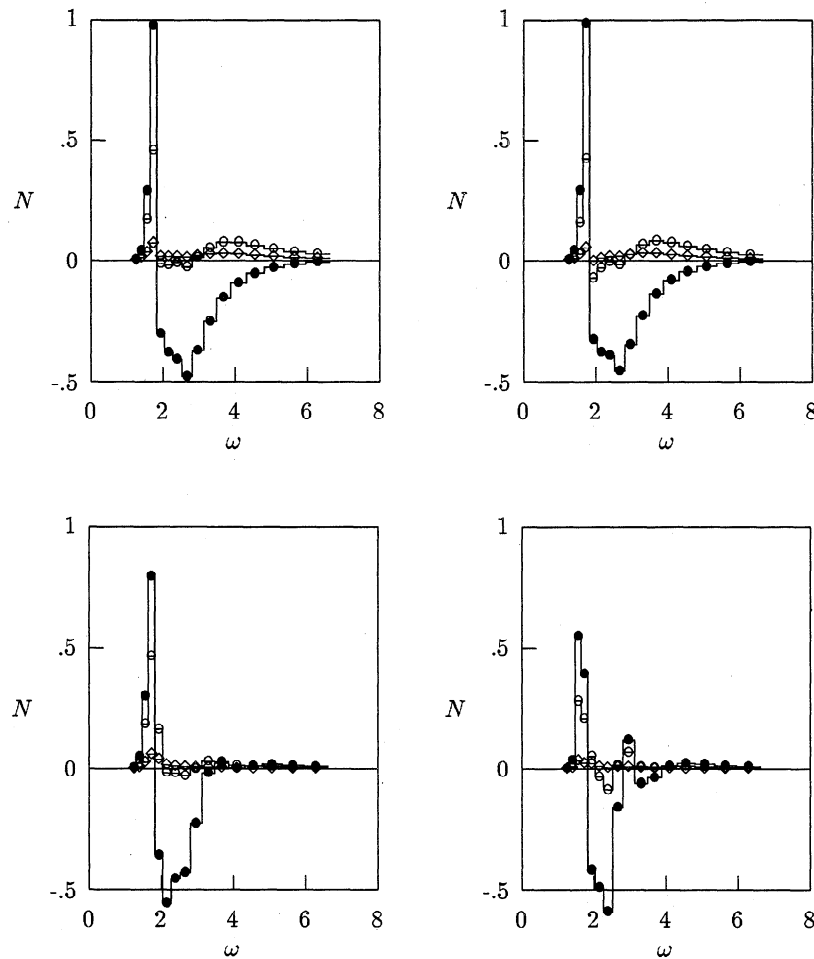
Figure 6. Differential transfer for various renormalized discrete-interaction truncations of Table 6. Computed for JONSWAP spectrum (6) with peak frequency .3 Hz. Displayed for angle bands 1 and 12 (bullets), 2 and 11 (open circles), and 3 and 10 (diamonds). Upper left panel is full hybrid transfer. Upper right panel is AT2 truncation. Lower left panel is DI5 truncation. Lower right panel is DI1 truncation. Units are 10^{-6} m^2 s and rad s^{-1}.

The reduced-integration truncations RI1 and RI2 do not appear to have much to recommend them. The efficiency of the RI1 truncation is comparable with that of the AT2 truncation, but its accuracy, as measured by both P_N and P_A, is not as good as that of the considerably more efficient DI5 truncation. The RI2 truncation is more accurate than the RI1 truncation, but is little more than half as accurate as and three times less efficient than the AT2 truncation. This truncation employs more than a third of the coefficients of the full hybrid sum.

While the reduced-integration approximation would seem to have a firm foundation in the singular structure of the interaction coefficient σ_N appearing in the Boltzmann integral (3), one is forced to conclude that (a) there are many interactions in the neighborhood of this singularity that are not important to the overall transfer and, conversely, (b) some interactions that are important to this transfer lie outside of this neighborhood. These considerations ultimately limit the efficiency and accuracy, respectively, of this approximation.

While the efficiency of the original approximation, which embeds the Boltzmann integration in the model integration, can certainly be improved relative to the RI1 and RI2 hybrid implementations by employing very coarse Boltzmann resolu-

tion, this improvement in efficiency must inevitably seriously degrade the accuracy of the approximation. Note also that the ratio between P_A and P_N is considerably larger for the RI2 truncation than it is for the other truncations of Table 6. This suggests that the cumulative effects of truncation error may be particularly troublesome for model integrations employing the reduced-integration approximation.

4. Systematic Truncation of the Hybrid Representation

The success of the hybrid generalizations of the discrete-interaction approximation, obtained by defining new selection algorithms generalizing the original algorithm, suggests that there may well exist other more effective truncations of hybrid transfer. We next outline a method for systematically identifying such truncations.

We start by recognizing that, if we fix the number of hybrid coefficients in the truncated sum, the optimum choice of these coefficients, with minimum variance between the renormalized truncated estimate of the nonlinear transfer and the full nonlinear transfer, is spectrum-dependent. The optimum choice for one input spectrum cannot be expected to be op-

Figure 7. Homogeneous evolution of the angle-integrated action spectrum for various renormalized discrete-interaction truncations of Table 6. Upper left panel employs full hybrid transfer. Upper right panel employs AT2 truncation. Lower left panel employs DI5 truncation. Lower right panel employs DI1 truncation. Piecewise-constant spectra are numerical integrations, shown at hourly intervals. Continuous spectrum is Phillips equilibrium spectrum. Units are m^2 s^2 and rad s^{-1}.

timium for other input spectra as well. Thus, if we seek an accurate and efficient representation of this transfer over a full range of input spectra, such as might represent various stages in the SNLdV model computations, we must either adapt the choice of coefficients to the developing spectrum or focus this optimization on a range of input spectra rather than on a single input spectrum. The first of these measures promises relatively high accuracy and efficiency, but requires some sort of adaptive selection of hybrid coefficient file(s) integral to the wave-model computation. Depending on its complexity, this selection may adversely impact the efficiency of this computation. More important, the optimization of coefficients for given spectral input (to define the coefficient files available to this selection) is computationally problematic if the focus is on individual coefficients rather than on groups of coefficients. Accordingly, while both measures can be accomodated by the SNLdV inverse-modeling program and while the second measure promises less accuracy and efficiency, we here discuss only this second measure.

One way to focus the choice of coefficients on a range of input spectra, suggested by STHHB, is first to determine the optimum coefficients specific to multiple input spectra and then to merge these coefficients. The present scheme essentially merges these coefficients up front by identifying and truncating the hybrid coefficients by group. It also explicitly incorporates multiple input spectra into the selection process.

We identify as belonging to a single group of hybrid coefficents all coefficients for which the bin difference between corresponding pairs of wave-number bins are the same and

for which the bin difference between corresponding pairs of angle bins are either the same or the same but of opposite sign (mirror images). Because the wave-number bins are geometrically spaced, this specification gathers together in single groups all interactions occupying the same or mirror-image positions in the interaction diagram (relative to the vector variable $\hat{k} \equiv k_1 + k_2 = k_3 + k_4$). (Note that other equivalent exchange interactions are already built into the computation by HH's symmetrization of the Boltzmann integral.) In terms of this grouping, the discrete-interaction truncations DI1, DI2, and DI3 contain 1, 2, and 3 coefficient groups, respectively.

This grouping accomplishes two things. First it guarantees that, if a set of groups of hybrid coefficients gives a good representation of the nonlinear transfer for a given input spectrum, it will give a good representation for other input spectra of similar shape but with different peak frequency. Second it guarantees that, if it gives a good representation of the nonlinear transfer for a given directional distribution, it will give a good representation for the mirror-image directional distribution. Both of these characteristics are necessary in order that the representation be appropriate to a full range of input spectra.

4.1. Scaling Considerations

To more explicitly develop this result, we note that the JONSWAP parameterization (6), with suitably chosen power laws expressing the dependence of the JONSWAP parameters

ϵ, Ω, γ, and σ on dimensionless fetch, provides a good description of the fetch-limited spectrum that develops at various distance from a long straight coast in response to a constant offshore wind [*Hasselmann et al.*, 1973]. While such a parameterization by no means encompasses all possible realizations of the spectrum, it does provide a reasonable range of input on which to base a standard choice of coefficients. Moreover, we can significantly enlarge this range by ignoring the dependence of the JONSWAP parameters on dimensionless fetch and regarding these parameters as independently variable.

We note first that in deep water the JONSWAP relation (6) may be written in the form,

$$A(k) = g^4 \Omega^{-9} \epsilon \chi(\gamma, \sigma, \kappa) \psi(\vartheta),$$

where χ is the dimensionless function,

$$\chi(\gamma, \sigma, \kappa) \equiv \frac{1}{2} \kappa^{-\frac{9}{2}} e^{-\frac{5}{4}\kappa^{-2}} \gamma^{e^{-\frac{1}{2\sigma^2}(\sqrt{\kappa}-1)^2}},$$

and κ is the dimensionless vector wave number,

$$\kappa \equiv g\Omega^{-2} k,$$

with magnitude κ and direction ϑ.

It follows that for JONSWAP input the Boltzmann integral (3) is of the form

$$N(k) = g^4 \Omega^{-8} \epsilon^3 I_N(\gamma, \sigma, \kappa, \vartheta), \tag{7}$$

where I_N is a dimensionless integral over the dimensionless variables κ_1, κ_2, and κ_3 of a dimensionless integrand derived from the integrand of (3). This integrand is a product of a dimensionless version of the interaction coefficient (a function of κ_1, κ_2, κ_3, and κ), the delta function factor,

$$\delta(\kappa_1 + \kappa_2 - \kappa_3 - \kappa)\delta(\sqrt{\kappa_1} + \sqrt{\kappa_2} - \sqrt{\kappa_3} - \sqrt{\kappa}),$$

and the spectral product factor,

$$\chi(\gamma, \sigma, \kappa_1)\psi(\vartheta_1)\chi(\gamma, \sigma, \kappa_2)\psi(\vartheta_2)$$
$$\big(\chi(\gamma, \sigma, \kappa_3)\psi(\vartheta_3) + \chi(\gamma, \sigma, \kappa)\psi(\vartheta)\big)$$
$$- \big(\chi(\gamma, \sigma, \kappa_1)\psi(\vartheta_1) + \chi(\gamma, \sigma, \kappa_2)\psi(\vartheta_2)\big)$$
$$\chi(\gamma, \sigma, \kappa_3)\psi(\vartheta_3)\chi(\gamma, \sigma, \kappa)\psi(\vartheta).$$

Suppose that for given ϵ, γ, σ, and $\psi(\vartheta)$ and for given peak frequency Ω equal to one of the bin frequencies of the spectral representation, a given set of hybrid coefficient groups optimizes the fit between the resulting estimate for the renormalized truncated nonlinear transfer and the full transfer (7). It follows that this set of coefficient groups also optimizes the corresponding estimate for the dimensionless integral I_N. Both estimates are a sum of renormalized contributions from each coefficient of each coefficient group. But not all coefficients from each coefficient group are important to the resulting sum. Within each group these coefficients may be ordered by wave number (of any of the interacting wave components). For given Ω, the contribution of the coefficients within a group is typically largest for some intermediate wave number in this ordering and becomes progressively less important at both lower and higher wave numbers.

Now suppose that the peak frequency of the input spectrum is displaced to some other bin frequency of the spectral representation. The principal effect of this displacement is to shift the center of gravity of the important hybrid coefficients within each coefficient group to a different wave number. This shift involves exactly the same displacement of wave-number bins that is involved in displacing the peak frequency bin and yields the same dimensionless vector wave numbers κ_1, κ_2, κ_3, and κ for the important coefficients as before. Furthermore, ignoring end effects at low and high wave number and the slightly incomplete convergence of the estimates for the different hybrid coefficients, the resulting contributions to the dimensionless integral I_N of the more important coefficients for the new peak frequency are essentially identical to the contributions to this integral of the more important coefficients for the old peak frequency. Thus the resulting estimate for I_N is the same, and the resulting estimate for the renormalized truncated nonlinear transfer for the new peak frequency is again optimal or very nearly so.

What if the displacement of the peak frequency is to a frequency intermediate between two bin frequencies? In this case, the important coefficients within each coefficient group are a combination of the important coefficients for the two bin frequencies. While the resulting performance of the truncation is more difficult to assess, we would expect this performance to again remain optimal or very nearly so.

What about the dependence on the JONSWAP parameters ϵ, γ, and σ and on the directional distribution $\psi(\vartheta)$? The integral I_N does not in fact depend on the Phillips parameter ϵ. Therefore, the optimum choice of coefficient groups is independent of this parameter. This integral does depend, however, on γ and σ. Thus one might anticipate that this choice might change with these parameters. JONSWAP finds some scatter but little detectable trend in the dependence of the parameters γ and σ on dimensionless fetch [*Hasselmann et al.*, 1973]. To the extent that these parameters may be considered constant, it follows that the optimum choice of coefficient groups for the nominal spectrum (6) should provide a reasonable standard for computing nonlinear transfer in the fetch-limited case. In fact, as demonstrated below, we find that this choice also gives reasonable results for values of γ that differ significantly from the nominal values characteristic of the fetch-limited case. We speculate that these results will also be relatively insensitive to the value of σ and to the choice of directional distribution $\psi(\vartheta)$.

To further smooth the dependence of the optimum choice of coefficient groups on the JONSWAP parameters, we allow this choice to incorporate multiple input spectra, minimizing the mean proportional variance for these inputs. While the input spectra for the optimizations reported below have in fact varied only the peak frequency Ω, they could also have varied γ, σ, and $\psi(\vartheta)$.

It should be remembered that we are dealing here with rather severe and greatly amplified (renormalized) truncations of the nonlinear transfer. The reason for considering such truncations is almost entirely practical. We need them to jump start a computation that would otherwise be too consuming of computer resources. These truncations are most probably not sufficiently accurate to carry the inverse-modeling program of SNLdV through to completion, but they are useful in getting it started. Furthermore, they grade naturally into less efficient but more accurate hybrid truncations, with renormalization factors of order 1, that should make it possible to complete this program.

4.2. Computational Details

An important practical advantage of focusing the selection process on groups of coefficients rather than on individual coefficients is that the overhead of this process can be greatly

reduced. This selection is aided by a simple transfer-based calculus arising from the specification (2). Let $N^{(j)}$ be the contribution to the nonlinear transfer from the jth coefficient group and let

$$C_{jk} \equiv \int d^2k \, N^{(j)} N^{(k)} \bigg/ \int d^2k \, N^2_{ref}$$

and

$$D_j \equiv \int d^2k \, N^{(j)} N_{ref} \bigg/ \int d^2k \, N^2_{ref}.$$

It follows that

$$P_N = 1 - \left(\sum_j D_j\right)^2 \bigg/ \left(\sum_{jk} C_{jk}\right)$$

and

$$\mu_N = \left(\sum_j D_j\right) \bigg/ \left(\sum_{jk} C_{jk}\right), \tag{8}$$

with a similar result for the case of optimizations that incorporate multiple input spectra. By precomputing the cross products C_{jk} and D_j, these relations allow a relatively rapid evaluation of P_N and μ_N for various combinations of coefficient groups.

Our scheme involves three distinct steps. (a) To limit the computational burden of the second and third of these steps, we start from an initial truncation of the full extrapolated case-4 hybrid representation. In what follows we have chosen two such truncations as starting points, specifically the interaction-angle truncations AT2 and AT3. (b) To further limit the computational burden of the third step, we next attempt to eliminate all remaining coefficient groups, the contribution of which is in some sense relatively small. This initial culling of groups employs one of two criteria. With the first criterion, groups are eliminated if $C_{jj} < C_{crit}$, with the second, if $|D_j| < D_{crit}$, where C_{crit} and D_{crit} are appropriate thresholds. This culling defines the pool of coefficient groups available to the subsequent optimization. (c) The proportional variance of various combinations of coefficient groups remaining in this pool is explicitly evaluated from (8) and an optimum combination of coefficient groups identified. This optimization is performed recursively in blocks of small dimension, starting from either the null representation (no coefficient groups) or from the discrete-interaction representation DI3 (three coefficent groups).

This scheme is limited primarily by the possible dimension of the recursive blocks. Let q be this dimension (block size), and let Q be the number of coefficient groups that survive the initial culling. Then the number of possible combinations of groups in the first recursive block of coefficient groups is given by $Q!/(q!(Q-q)!)$. It follows that the required computation time is a strong function of both Q and q. In practice, an initial culling to 153 coefficient groups (such as employed in most of the optimizations depicted in Figure 8) restricts the block size q to 5 or less. To enable a block size of 10, the hybrid coefficients must be culled to something like 41 coefficient groups. This level of culling significantly restricts the pool of coefficient groups available to the optimization. (For example, only two of the three coefficient groups defining the DI3 truncation survive this culling.)

4.3. Results

The resulting truncation of the AT2 and AT3 coefficient files for nominal spectral resolution, using multiple spectral inputs, is described by Figure 8. Shown in this figure are the transfer-based proportional variance P_N and renormalization μ_N for a series of truncations (cullings and subsequent recursive optimizations) of these files. These variables are shown as functions of the number of hybrid coefficients n_H surviving the truncations. Also included in the figure are the renormalized truncations of Table 6.

Two series of cullings are shown, those employing the first threshold criterion described above and those employing the second. In each case, the dependence of P_N and μ_N on n_H is generated from the parent truncations AT2 and AT3 by incrementally raising the corresponding thresholds and allowing fewer and fewer coefficient groups to survive. Note that considerable initial culling of the parent truncations is possible without appreciably degrading the resulting representation. Indeed, using the first criterion, both truncations may be trimmed to essentially one third their original size, with little change in P_N.

Additional modest culling of both parent truncations results in an intermediate truncation with 153 coefficient groups (4372 coefficients), with $P_N = .0132$ and $\mu = 1.19$, where the AT2 and AT3 branches of the dependence of P_N and μ_N on n_H have coalesced. This truncation defines the coefficient pool available to and common end point for subsequent recursive optimizations with block sizes of 1, 2, 3, 4, and 5 coefficient groups. Further culling of this intermediate truncation to 119, 96, 83, and 72 coefficient groups defines suitable coefficient pools for optimizations with block sizes of 6, 7, 8, and 9 coefficient groups. These optimizations define the renormalized combinations of coefficient groups that provide optimum representations for the nonlinear transfer. They start from either the null truncation or the DI3 truncation and end with the full intermediate truncation.

Analysis of these optimizations suggests the following. (a) Recursive block optimization dramatically lowers the proportional variance relative to hybrid implementations of the discrete-interaction approximation DI1 through DI6 and to culled AT2 and AT3 truncations of similar dimension. (b) Different optimizations result in similar levels for P_N and μ_N, with some tendency for larger block size to result in smaller P_N. (c) For small numbers of coefficients, the optimizations starting from the null truncation are better than those starting from the DI3 truncation. (d) For somewhat larger numbers of coefficients, the optimizations starting from the DI3 truncation are better than those starting from the null truncation. (e) There is not much of a tendency for distinct optimizations to share the same coefficients. There are apparently many distinct combinations of coefficients that result in improved representations for the nonlinear transfer. In particular only one of the three coefficient groups that comprise the DI3 truncation is typically picked up by optimizations starting from the null group, and it is typically the last of the DI3 groups. (f) The coefficient groups identified by a particular optimization appear unrelated to one another. They come, not from single localized regions of the interaction diagram, but are spread throughout this diagram. The function of additional coefficient groups appears to be to complement rather than to augment. It is unlikely that one can summarize the selection of these groups in terms of a simple algorithm such as characterizes the discrete-interaction and reduced-integeration approximations. (g) In contrast to the culled AT2 and AT3 truncations, which are characterized by modest renormalizations of order 1 to 2, these optimizations

Figure 8. Optimum transfer-based proportional variance P_N and renormalization μ_N for various truncations of hybrid transfer as a function of the number of hybrid coefficients retained n_H: large bullets, selected truncations including the DI1-DI6, RI1, RI2, AT2, AT3, CL1, OP5, and OP9 truncations; medium bullets, culled truncations of type 1, starting from AT2 and AT3; and small bullets, culled truncations of type 2, starting from AT2 and AT3. Remaining curves depict recursive optimizations of block size 1, 2, 3, 4, and 5, starting from the DI3 truncation or from the null truncation. All truncations renormalized to minimize the mean proportional variance for JONSWAP input spectrum (6) with peak frequencies .3, .4, .5, .6, and .7 Hz.

typically require substantial renormalizations of order 7 or more.

Because the intermediate truncation that defines the coefficient pool for a block-size-5 optimization does not contains all coefficient groups, we cannot be absolutely certain that this optimization identifies the very best combination of 5 coefficient groups (among all possible coefficient groups). Similarly, because the intermediate truncation that defines the pool for a block-size-10 optimization is considerably smaller, there is even less justification for claiming that this optimization identifies the best combination of 10 coefficient groups. (The second iteration of the block-size-5 computation does not identify this combination either, because it includes only combinations that contain the same first five coefficient groups.)

But, while it is clear that this scheme cannot practically look at all possible combinations of coefficient groups contained in the AT2 and AT3 truncations, it also seems likely that this failing is of little consequence. What is important is that, for the level of efficiency defined by the product of the block size and the number of recursive iterations,

these optimizations are probably about as accurate as can be constructed. Moreover, it makes little practical difference whether our scheme identifies the most accurate combination of coefficient groups or simply identifies a combination that is almost as accurate. Indeed, our scheme identifies many such combinations of coefficient groups.

We next focus on three particular combinations of coefficient groups defined by this optimization scheme, the truncations CL1, OP5, and OP9 of Figure 8. CL1 is a culled truncation of type 1 with 742 coefficient groups. OP5 is a recursive block optimization of block size 5 with 33 coefficient groups, selected from a culled truncation of type 1 with 153 coefficient groups, starting from the DI3 truncation. OP9 is a single-iteration block optimization of block size 9 with nine coefficient groups, selected from a culled truncation of type 1 with 72 coefficient groups, starting from the null truncation.

Table 7 and Figures 9 through 11 compare the transfer estimates for these truncations with those of the full extrapolated case-4 hybrid sum, computed for the JONSWAP input spectrum (6). Table 7 shows the relative insensitivity to peak

Table 7. Transfer-Based Truncation Error for Three
Systematic Truncations

Case	n_H	μ_N	$P_N^{(.3)}$	$P_N^{(.4)}$	$P_N^{(.5)}$	$P_N^{(.6)}$	$P_N^{(.7)}$
CL1	19,392	1.01	.0002	.0001	.0002	.0002	.0002
			.0003	.0003	.0003	.0003	.0003
OP5	942	8.45	.0025	.0009	.0015	.0014	.0019
			.0040	.0043	.0042	.0068	.0105
OP9	256	11.7	.0061	.0029	.0028	.0052	.0073
			.0122	.0122	.0135	.0180	.0400

Here n_H is the number of hybrid coefficients in the truncation, μ_N is the renormalization factor, and the P_N are the proportional variances between the resulting transfer and full hybrid transfer for JONSWAP spectrum (6) with peak frequency .3, .4, .5, .6, and .7 Hz and peak enhancement 3.3 (upper entry) and 1.0 (lower entry).

frequency Ω and peak enhancement γ of the resulting proportional variance P_N between the renormalized truncated estimates and full transfer. This table shows some deterioration in the representations provided by the optimized truncations OP5 and OP9 for extreme values of the peak frequency and for the lower choice of peak enhancement, but this deterioration is modest. The lower value of peak enhancement (1.0) typically gives only a fourfold increase in proportional variance

relative to the higher value (3.3), even though the choice of coefficients employs only input spectra with the higher value.

Figures 9 and 10 show the corresponding angle-integrated transfer for these truncations (and for the full hybrid sum) for three values of peak frequency (.3, .4, and .5 Hz) and two values of peak enhancement (3.3 and 1.0). Figure 10 shows the differential transfer for a single value of peak frequency (.3 Hz) and peak enhancement (3.3). These figures should be compared with corresponding figures from the previous section. In all cases, the culled transfer CL1 is almost indistinguishable from full transfer and the match between the optimized estimates OP5 and OP9 and full transfer is considerably more realistic than that of corresponding hybrid implementations of the discrete-interaction approximation.

Table 8 shows the corresponding mean proportional variance between the homogeneous evolution of the spectrum for these truncations and this evolution for full hybrid transfer, computed both for the renormalized coefficients of Table 7 and for these coefficients rerenormalized to minimize this proportional variance (over the first 4 hours of this evolution). This table reveals a level of proportional variance that is somewhat larger than that of Table 7. Rerenormalization does little to change the level of this proportional variance. The OP5 truncation maintains a reasonably uniform level of
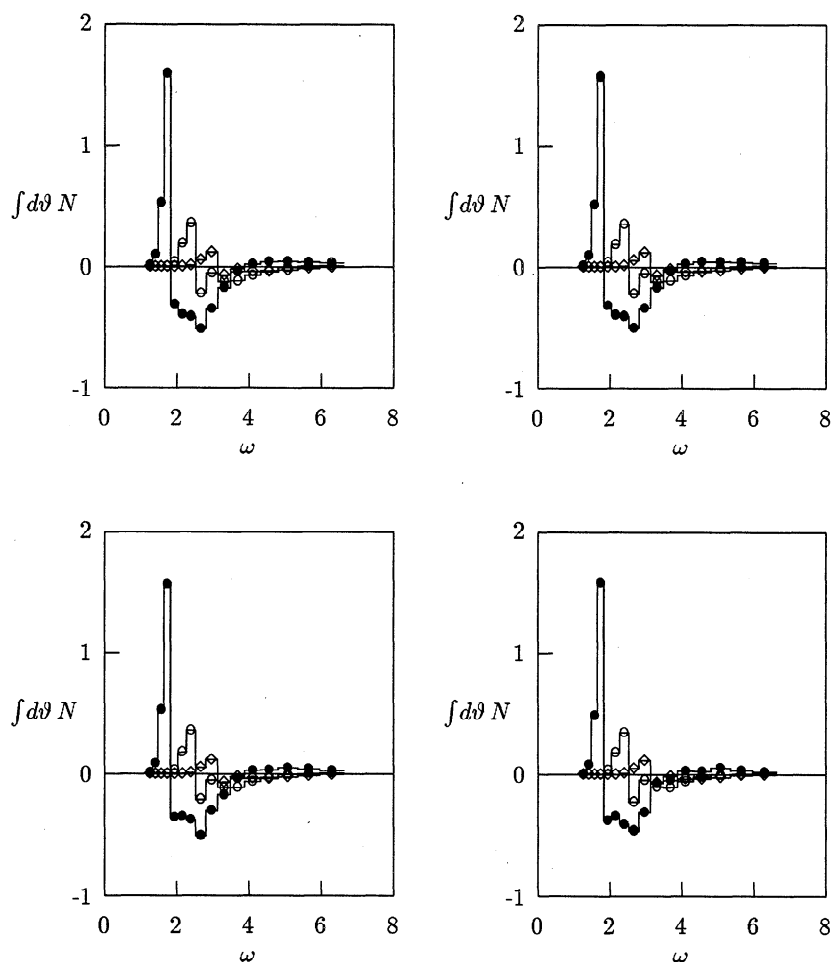

Figure 9. Angle-integrated transfer for various renormalized hybrid truncations of Figure 8. Computed for JONSWAP spectrum (6) with peak frequency .3 Hz (bullets), .4 Hz (open circles), and .5 Hz (diamonds). Upper left panel is full hybrid transfer. Upper right panel is CL1 truncation. Lower left panel is OP5 truncation. Lower right panel is OP9 truncation. Units are 10^{-6} m^2 s and rad s^{-1}.
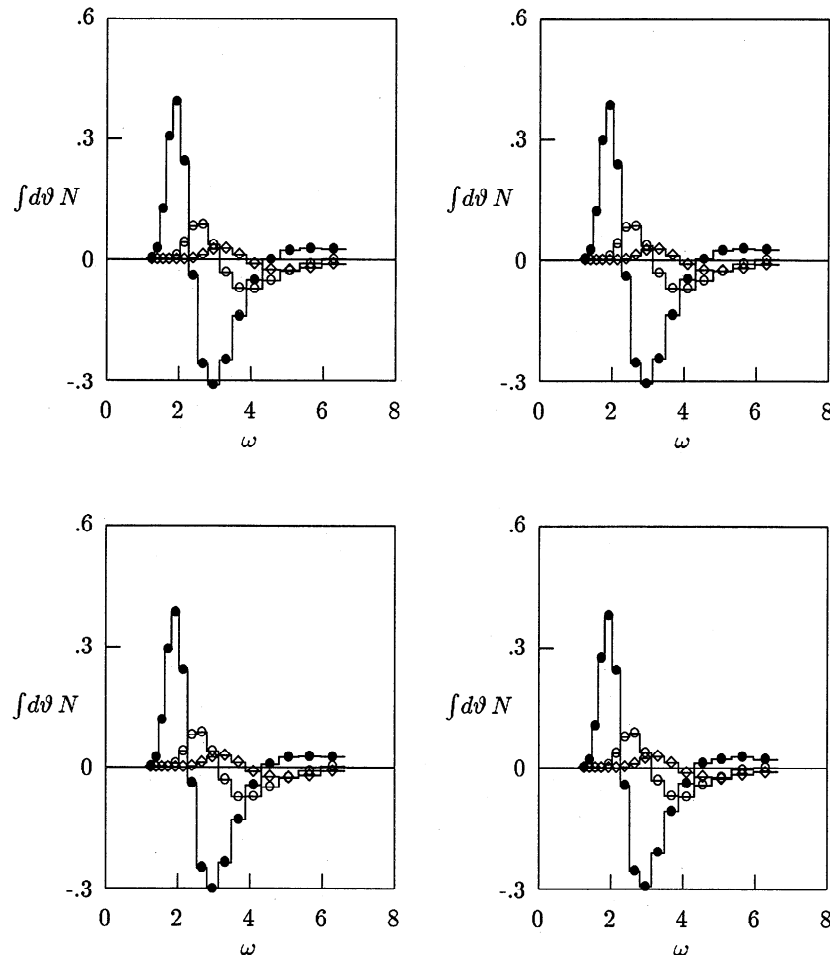
Figure 10. Angle-integrated transfer for various renormalized hybrid truncations of Figure 8. JONSWAP input spectra as in Figure 9, except peak enhancement is 1.0. Upper left panel is full hybrid transfer. Upper right panel is CL1 truncation. Lower left panel is OP5 truncation. Lower right panel is OP9 truncation. Units are 10^{-6} m^2 s and rad s^{-1}.

proportional variance throughout the evolution. The OP9 truncation gives a proportional variance that is significantly larger at the start of the evolution than it is at the end. The proportional variance for the fourth hour of the OP9 evolution is significantly smaller than for the fourth hour of the OP5 evolution. The implications of this result are not entirely clear. Probably, the lack of better OP9 performance exhibited in this figure at the start of the evolution is related to the deterioration in OP9 performance exhibited in Table 7 at the highest peak frequency. It may be the case that in an extended integration, where the peak frequency is mostly at lower levels, the OP9 truncation may perform very well indeed. It may also be the case that inclusion of the DI3 coefficient groups (as in the OP5 truncation) is necessary to ensure a reasonably consistent performance over the multiple stages of this evolution. On the other hand, this relative consistency may simply result from the fact that this truncation is longer.

Figure 12 shows the corresponding evolution of the angle-integrated spectrum for these truncations (and for full hybrid transfer), using the renormalized coefficients of Table 7. Clearly, the CL1 truncation is again almost indistinguishable from, and the OP5 and OP9 evolutions bear considerable resemblance to, that of the full hybrid sum. These evolutions exhibit a strong overshoot and a Phillips tail, the level of which is much improved relative to the evolutions displayed in the previous section for various hybrid implementations of the discrete-interaction approximation. Coincidentally, as evidenced by Figure 11, the angular dependence of the nonlinear transfer in this tail is also much improved.

4.4. Extension to Higher Spectral Resolution

As previously remarked, the present results have been developed for a relatively coarse nominal spectral resolution of 16 prognostic and 4 diagnostic wave-number bands and 12 angle bands. While it is clear that this nominal resolution provides a suitable starting point for the SNLdV inverse-modeling program, it is also clear that to contain the representation error, this program must eventually progress beyond nominal spectral resolution to at least double spectral resolution. (Fortunately, as previously discussed, it also appears that double resolution ought to be sufficient to realize an extrapolated representation error of a few percent.)

Whether in relation to this inverse modeling or in relation to other wave-modeling applications where nominal resolution may be too coarse, it is of interest to inquire to what extent the recursive-optimization scheme described in this section can be extended to higher spectral resolution.
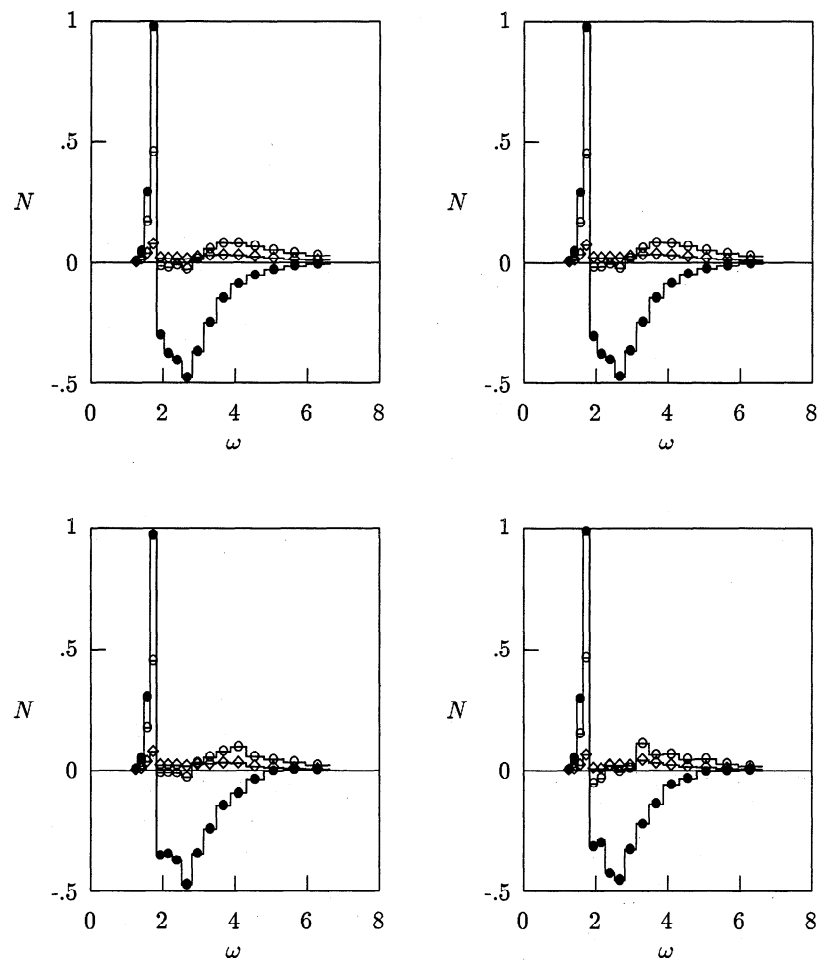
Figure 11. Differential transfer for various renormalized hybrid truncations of Figure 8. Computed for JONSWAP spectrum (6) with peak frequency .3 Hz. Displayed for angle bands 1 and 12 (bullets), 2 and 11 (open circles), and 3 and 10 (diamonds). Upper left panel is full hybrid transfer. Upper right panel is CL1 truncation. Lower left panel is OP5 truncation. Lower right panel is OP9 truncation. Units are 10^{-6} m^2 s and rad s^{-1}.

The basic difficulty in going to higher spectral resolution is that (a) the number of coefficient groups in the full hybrid sum is a strong increasing function of this resolution, and (b) the computer time required for the final selection of coefficient groups is a strong increasing function of the size of the culled pool of coefficient groups from which these coefficient groups are selected. This implies that, for higher resolution, in order

Table 8. Evolution-Based Truncation Error for Three Systematic Truncations

Case	n_H	μ_A	$P_A^{(1)}$	$P_A^{(2)}$	$P_A^{(3)}$	$P_A^{(4)}$
CL1	19,392	1.01	.0002	.0003	.0001	.0001
		1.00	.0002	.0002	.0001	.0001
OP5	942	8.45	.0173	.0259	.0298	.0336
		10.1	.0189	.0198	.0203	.0200
OP9	256	11.7	.1295	.0856	.0143	.0070
		9.62	.1321	.0358	.0083	.0091

Here n_H is the number of hybrid coefficients in the truncation, μ_A is the optimizing renormalization factor (upper entry) or rerenormalization factor (lower entry), and the P_A are the proportional variances relative to a computation employing full hybrid transfer for hours 1, 2, 3, and 4 of the corresponding homogeneous evolution of the action spectrum.

to contain the computer time of the final selection, one must initially cull a larger proportion of coefficient groups and/or one must lower the block size for this selection.

With double resolution there are initially 15 times as many base coefficients as there are for nominal resolution. That translates to 8 times as many coefficient groups. While this may seem a large factor to overcome, our expectation is that an extension of the recursive-optimization procedure~to double spectral resolution is quite feasible and will yield some useful truncations analogous to those for nominal spectral resolution. Similarly, while we are less optimistic about extending these results to quadruple spectral resolution, we do not rule out the possibility that such an extension may also yield some useful truncations.

4.5. Possible Implications for Adaptive Modeling

One of the features of the hybrid representation for nonlinear transfer that we have not fully exploited is the open-ended character of the hybrid coefficient files at low and high wave number. Our model employs fixed frequency limits, so the only reason to extend these files is to allow for nonlinear interactions with a fixed number of diagnostic wave components. But instead of fixing these limits, this model could

Figure 12. Homogeneous evolution of the angle-integrated action spectrum for various renormalized hybrid truncations of Figure 8. Upper left panel is full hybrid transfer. Upper right panel is CL1 truncation. Lower left panel is OP5 truncation. Lower right panel is OP9 truncation. Piecewise-constant spectra are numerical integrations, shown at hourly intervals. Continuous spectrum is Phillips equilibrium spectrum. Units are $m^2 s^2$ and rad s^{-1}.

allow them to vary in geometric increments as functions of space and time, with values adaptively set by the local spectrum. While such an adaptive model would complicate the integration of the model equations, it might well result in a more efficient integration. The point is that the form of the hybrid sum would appear to be well suited to such an adaptive model. By extending the coefficient file to both lower and higher wave number, it should be possible to create a single coefficient file, contiguous subfiles of which would provide the necessary hybrid coefficients for a variety of adaptively set limits.

This possibility is particularly intriguing in the case of a deep-water computation. Here it may be possible to further exploit deep-water ($\hat{k}$) scaling to replace this coefficient file with a considerably shorter coefficient file that essentially collapses all of the coefficients belonging to a single coefficient group into two dimensionless coefficients. Corresponding to this collapse, the hybrid sum (5) would gain a second inner sum over the difference between the wave-number bin of a prognostic or diagnostic wave component and the wave-number bin of some standardizing feature such as the spectral peak. It might further be possible to combine these inner sums into a single vectorizable sum of large dimension, providing an alternative to the vectorization of the nonlinear-transfer computation on the spatial grid points of the model, which alternative could prove advantageous on a vectorizing platform.

5. Summary and Conclusions

In this first of a series of papers describing a fully nonlinear regional wave model for the Bight of Abaco, we have discussed primarily the numerical errors associated with the model's hybrid representation for the nonlinear transfer from wave-wave interactions, extending a number of results previously reported by STHHB, and we have developed several techniques for accurately and efficiently truncating (and renormalizing) this representation. The following summary and conclusions should be viewed in the broader context of the SNLdV inverse-modeling program that has prompted the development of this model.

1. We have doubled the resolution of the STHHB Boltzmann integration to 648 wave-number bands and 648 angle bands, computing the hybrid coefficients for a nominal spectral resolution of 16 prognostic and four diagnostic wave-number bands and 12 angle bands and for a double spectral resolution of 32 prognostic and eight diagnostic wave-number bands and 24 angle bands. Linear extrapolation of the results of these integrations to infinite integration resolution suggests an integration error of less than 1% (convergence to better than 1%) for both spectral resolutions.

2. Comparison of less highly convergent estimates for nominal, double, and quadruple spectral resolution suggests a corresponding representation error of 13% for nominal resolution and 4% for double resolution. We speculate, however, that

the impact of this error on the the inverse modeling can be reduced to the level of a few percent by extrapolating the results of the inverse modeling for nominal and double spectral resolution to the case of infinite spectral resolution (in much the same way that the results of the Boltzmann integration have been extrapolated to the case of infinite integration resolution). While it is thus be necessary to do some inverse modeling at double resolution to contain this error, this will only be necessary during the later stages of the inverse modeling.

3. The double-resolution hybrid computation better resolves the negative mid-frequency lobe of the angle-integrated nonlinear transfer for JONSWAP input than does the nominal-resolution computation, confirming the double peak previously reported by STHHB, but suggesting that this double peak is more appropriately characterized as a superposition of two peaks, a very sharp peak at the low-frequency margin of this mid-frequency lobe and a broad peak extending across the entire lobe.

4. The use of diagnostic bands is an effective way to account for wave-wave interactions involving high-wave-number wave components external to the prognostic range of the spectrum. It does not take too many diagnostic bands to reduce the error associated with a further lack of these bands to acceptable limits. This error depends on the peak frequency. For nominal spectral resolution and an upper prognostic frequency of 1 Hz, four diagnostic bands are sufficient to reduce this error to 2% or less for peak frequencies less than or equal to .5 Hz.

5. By splitting the nonlinear-transfer computation into two computations, the first a deep-water computation encompassing interactions among wave components, the wave numbers for which all exceed some depth-dependent threshold, and the second a depth-dependent computation encompassing all other interactions, it is possible to extend the hybrid scheme to finite depth without greatly increasing the burden of the computation. We discretize the depth in five increments and employ a single set of deep-water coefficients and five correspondingly restricted sets of finite-depth coefficients.

6. Selective truncation and renormalization of the hybrid sum yields hybrid equivalents and generalizations for several existing approximations to the full nonlinear transfer. In particular, we have evaluated a number of truncations with a selection algorithm similar to that of the discrete-interaction approximation. The simplest of these truncations, the DI1 truncation, involves the same number of base interactions as and gives transfer estimates comparable to (but not identical to) the discrete-interaction approximation. The remaining discrete-interaction truncations relax the constraint that the incoming wavenumbers k_1 and k_2 be identical, giving progressively more accurate but less efficient hybrid representations for the nonlinear transfer.

7. We have also evaluated several hybrid truncations with selection algorithms similar to that of the reduced-integration approximation recently proposed by R. Q. Lin et al. (unpublished manuscript, 1996). These truncations are more accurate than but considerably less efficient than the discrete-interaction approximation. Their accuracy-efficiency product is an order of magnitude smaller than that of the interaction-angle truncations of Section 3 and more than 2 orders of magnitude smaller than that of the recursively optimized truncations of Section 4. In its original form, the reduced-integration approximation can realize greater efficiency by employing a very coarse Boltzmann integration resolution, but this coars-

ening can be expected to correspondingly degrade the accuracy of the approximation (relative to these hybrid truncations).

8. Finally, we have developed a systematic procedure for culling and recursively optimizing truncations of the hybrid sum and have employed this procedure to identify a sequence of optimized hybrid truncations of increasing accuracy, suitable for a staged implementation of the SNLdV inverse modeling of unknown action-balance source terms. This procedure represents an alternative to the merging of input-specific truncations suggested by STHHB to obtain a truncation with good performance over a range of spectral input. It achieves this performance up front by selecting coefficient groups (the interactions for which are scaled versions of one another or of mirror images of one another) rather than individual coefficients and by explicitly incorporating a range of input spectra into the culling and optimization procedures.

9. We start from the AT2 and AT3 interaction-angle truncations, culling these truncations by eliminating coefficient groups whose contribution to the nonlinear transfer is in some sense small. This culling defines the pool of coefficients available to the subsequent optimization. Starting from either the null truncation or some other truncation (typically the DI3 truncation), we build the optimized truncations by recursively adding optimizing coefficient groups in blocks of small dimension.

10. The culling and recursive-optimization stages of this procedure both take advantage of a simple calculus that attends a transfer-based specification of the error. This calculus determines the minimizing renormalization factor μ_N and corresponding minimum mean proportional variance P_N for a given truncation. It further determines these variables in terms of some basic integrals over coefficient groups and pairs of coefficient groups that need be computed only once in order to engage the recursive optimization, resulting in a very efficient selection of successive blocks of coefficient groups.

11. The resulting truncations include a number of truncations that essentially trade a factor of 10 decrease in efficiency for a factor of 10 increase in accuracy (relative to the discrete-interaction approximation). These truncations define a reasonable starting point for the SNLdV inverse modeling of unknown source terms and provide a reasonable alternative to the discrete-interaction approximation currently employed in many operational wave models. A second truncation, a culled version of the AT3 interaction-angle truncation, requires little or no renormalization, is almost as accurate as the full hybrid representation, and is only 600 times slower than the discrete-interaction approximation. While such a decrease in efficiency may not be acceptable for operational wave models, it is probably quite acceptable for the final stages of the SNLdV inverse modeling.

12. There is nothing truly unique about the recursive optimizations. They all typically require considerable renormalization. They do not show much tendency to share the same coefficient groups. They do not show much tendency to pick up the coefficient groups of the discrete-interaction truncations. While there are some systematic differences among these truncations, depending on block size and starting truncation, one recursive optimization of given dimension is about as good as another. In short, there are many seemingly unrelated choices of coefficient groups that result in much the same much-improved working representation for the nonlinear transfer (relative to the discrete-interaction approximation).

13. Extension of this recursive optimization procedure to higher spectral resolution is complicated primarily by the large increase in the number of coefficient groups that attends the increase in resolution. Accordingly, to contain the computer time required for the selection of optimizing groups, one must initially cull a larger proportion of these groups and/or one must lower the block size for this selection. Nonetheless, we anticipate that the extension of this procedure to double spectral resolution will be relatively straightforward and will yield some useful truncations suitable for both inverse wave modeling and operational wave modeling.

References

Beji, S., and J. A. Battjes, Experimental investigation of wave propagation over a bar, *Coastal Eng. J., 19*, 151–162, 1993.

Hasselmann, K., On the nonlinear energy transfer in a gravity-wave spectrum, 1, General Theory, *J. Fluid Mech., 12*, 481–500, 1962.

Hasselmann, K., On the nonlinear energy transfer in a gravity-wave spectrum, 2, Conservation theorems, wave-particle correspondence, irreversibility, *J. Fluid Mech., 15*, 273–281, 1963a.

Hasselmann, K., On the nonlinear energy transfer in a gravity-wave spectrum, 3, Computation of the energy flux and swell-sea interaction for a Neumann spectrum, *J. Fluid Mech., 15*, 385–398, 1963b.

Hasselmann, K., Weak interaction theory of ocean waves, *Basic Dev. Fluid Dyn., 2*, 117–182, 1968.

Hasselmann, K., et al., Measurements of wind-wave growth and decay during the Joint North Sea Wave Project (JONSWAP), *Dtsch. Hydrogr. Z., A8*(12), 1–95, 1973.

Hasselmann, K., On the spectral dissipation of ocean waves due to whitecapping, *Boundary Layer Meteorol., 6*, 107–127, 1974.

Hasselmann, S., and K. Hasselmann, A symmetrical method of computing the nonlinear transfer in a gravity wave spectrum, *Hamb. Geophys. Einzel-schr., A, 52*, 138 pp., 1981.

Hasselmann, S., K. Hasselmann, J. H. Allender, and T. P. Barnett, Computations and parameterizations of the nonlinear energy transfer in a gravity-wave spectrum, II, Parameterizations of the nonlinear energy transfer for application in wave models, *J. Phys. Oceanogr., 15*, 1378–1391, 1985.

Herterich, K., and K. Hasselmann, A similarity relation for the nonlinear energy transfer in a finite-depth gravity-wave spectrum, *J. Fluid Mech., 97*, 215–224, 1980.

Janssen, P. A. E. M., Quasi-linear theory of wind wave generation applied to wave forecasting, *J. Phys. Oceanogr., 21*, 1631–1642, 1991.

Komen, G. J., L. Cavaleri, M. Donelan, K. Hasselmann, S. Hasselmann, and P. A. E. M. Janssen, *Dynamics and Modeling of Ocean Waves*, 532 pp., Cambridge Univ. Press, New York, 1994.

Miles, J. W., On the generation of waves by shear flows, *J. Fluid Mech., 3*, 185–204, 1957.

Phillips, O. M., On the generation of waves by turbulent wind, *J. Fluid Mech., 2*, 417–445, 1957.

Resio, D. S., and W. Perrie, A numerical study of nonlinear emergy fluxes due to wave-wave interactions, 1, Methodology and basic results, *J. Fluid Mech., 223*, 603–629, 1991.

Snyder, R. L., W. Neu, R. B. Long, and W. J. P. de Voogt, A long-range program to parameterize the two-dimensional evolution of the surface-gravity-wave field, Nova Univ. Tech. Rep., *90-2*, 30 pp., Nova Southeastern Univ., Dania, Fla., 1990.

Snyder, R. L., L. M. Lawson, and R. B. Long, Inverse modeling of the action-balance equation, I, Source expansion and adjoint-model equations, *J. Phys. Oceanogr., 22*(12), 1540–1555, 1992.

Snyder, R. L., W. C. Thacker, K. Hasselmann, S. Hasselmann, and G. Barzel, Implementation of an efficient scheme for calculating nonlinear transfer from wave-wave interactions, *J. Geophys. Res., 98*(C8), 14,507–14,525, 1993.

Thacker, W. C., Some computational problems of oceanography, in *Proceedings Fourth International Conference on Finite Elements in Water Resources*, Springer-Verlag, New York, 1982.

Thacker, W. C., and R. B. Long, Fitting dynamics to data, *J. Geophys. Res., 93*(C2), 1227–1240, 1988.

Tolman, H. L., Effects of numerics on the physics in a third-generation wind-wave model, *J. Phys. Oceanogr., 22*, 1095–1111, 1992.

WAMDIG, The WAM model – A third generation ocean wave prediction model, *J. Phys. Oceanogr., 18*, 1775–1810, 1988.

Webb, D. J., Nonlinear transfer between sea waves, *Deep-Sea Res., 25*, 279–298, 1978.

R. B. Long and R. L. Snyder, Oceanographic Center, Nova Southeastern University, 8000 N. Ocean Drive, Dania, FL 33004. (e-mail: long@tiger.ocean.nova.edu; snyder@tiger.ocean.nova.edu)

W. L. Neu, Department of Aerospace and Ocean Engineering, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061. (e-mail: neu@neu.aoe.vt.edu)