

PIK Report

No. 68

A PARALLEL ALGORITHM
FOR THE DISCRETE ORTHOGONAL
WAVELET TRANSFORM

Markus Uhlmann



POTSDAM INSTITUTE
FOR
CLIMATE IMPACT RESEARCH (PIK)

Author:

Dr. Markus Uhlmann

Potsdam Institute for Climate Impact Research

P.O. Box 60 12 03, D-14412 Potsdam, Germany

Phone: +49-331-288-2687

Fax: +49-331-288-2695

E-mail: Markus.Uhlmann@pik-potsdam.de

Herausgeber:

Dr. F.-W. Gerstengarbe

Technische Ausführung:

U. Werner

POTSDAM-INSTITUT
FÜR KLIMAFOLGENFORSCHUNG
Telegrafenberg
Postfach 60 12 03, 14412 Potsdam
GERMANY

Tel.: +49 (331) 288-2500

Fax: +49 (331) 288-2600

E-mail-Adresse: pik-staff@pik-potsdam.de

POTSDAM, DEZEMBER 2000

Abstract

We recall the basics of discrete orthogonal wavelet bases and show how a fast algorithm for the transform of n -dimensional data can be constructed and implemented on distributed memory machines. For this purpose, we use a ‘slice’ representation of data across processors and restrict to the case of a power-of-two number of processors for simplicity. Some examples of the transform and filtering of two- and three-dimensional data are given. It is found that our parallel data-model leads to a satisfactory scalability of the algorithm.

1 Introduction

Wavelet analysis has proven useful in applications as different as image compression (e.g. Mallat, 1989) and meteorology (e.g. Torrence & Compo, 1998), to name but two examples. For nearly a decade, wavelet-based techniques have also been applied to various problems in fluid mechanics, either for (turbulent) data analysis (Meneveau, 1991; Do-Khac, Basdevant, Perrier & Dang-Tran, 1994; Brasseur & Wang, 1992; Fröhlich, 2000), data compression (Farge, Schneider & Kevlahan, 1998) or with the purpose of constructing numerical methods (e.g Fröhlich & Schneider, 1994). Our present motivation is of the first type. Within a turbulent flow, we wish to detect certain events which are associated to some length scale *and* appear localized in space (Klein, 1998). These are quite exactly the features which make wavelet functions so attractive: they offer a compromise between Fourier analysis and physical space analysis allowing for a scale-wise study of the signal at different points in space.

The present report is intended to document our method of analysis in detail. For a general introduction to the theory of wavelets – from a mathematical point of view – the reader is referred to the classical monographs by Meyer (1992) and Daubechies (1992). Here we have chosen to work with orthonormal bases of discrete wavelets. The main reason is the fact that the data volume is kept constant under these circumstances, which is very desirable in view of the large number of degrees of freedom of the turbulent flow fields (Meneveau, 1991). Moreover, the orthogonality and completeness of the basis allows us to perform filtering operations in wavelet space ensuring that the back-transformed field remains smooth. What is sometimes stated as a disadvantage of the discrete transform is the real-valued nature of the wavelet functions (and therefore of the coefficients) which leads to small-scale oscillations of the ensuing coefficients (Do-Khac *et al.*, 1994). On the other hand, it should be pointed out that the study of the modulus of complex wavelet coefficients (obtained via continuous transform) suppresses information about the sign of the signal.

The organization of this document is as follows. Section 2 introduces the discrete wavelet transform particularly for n -dimensional periodic data and using spline wavelets. In §3 we describe how the fast transform algorithm can be implemented on a distributed memory machine. Section 4 finally shows some examples of two- and three-dimensional transforms as well as measures of performance of the parallel data-model.

2 The discrete orthogonal wavelet decomposition

2.1 Basics

We wish to construct an orthonormal basis of the space of square integrable functions, consisting of translations and dilations of essentially one “mother” function in order to enable us to analyze a given function $f(x) \in L^2(\mathbb{R}) \equiv V$ with respect to space and scale. According to the theory of multi-resolution analysis (MRA, Mallat, 1989), we will work with two orthogonal functions:

(i) the scaling function ϕ , acting as a low-pass filter or averaging function, where

$$\int_{-\infty}^{+\infty} \phi(x) dx = 1 \quad (1)$$

holds;

(ii) the basic wavelet ψ , which is a band-pass filter or a “detail” function with a mean value of zero,

$$\int_{-\infty}^{+\infty} \psi(x) dx = 0 \quad . \quad (2)$$

We define the dyadic dilation (by a factor of 2^{-j}) and the translation (by s) of a function $g(x)$ as follows:

$$g_{j,s}(x) = 2^{j/2} \cdot g(2^j x - s) \quad . \quad (3)$$

Then, the necessary orthogonality conditions can be expressed as:

$$\langle \phi_{j,s}, \phi_{k,r} \rangle = \delta_{sr}, \quad (4a)$$

$$\langle \psi_{j,s}, \psi_{k,r} \rangle = \delta_{sr} \delta_{jk}, \quad (4b)$$

$$\langle \psi_{j,s}, \psi_{j,r} \rangle = 0, \quad (4c)$$

where $\langle a, b \rangle = \int_{-\infty}^{+\infty} ab dx$.

The orthogonal basis is constructed from both, ϕ and ψ , which leads to the following decomposition of our to-be-analyzed function:

$$f(x) = \sum_{s \in \mathbb{Z}} c_{j',s} \phi_{j',s}(x) + \sum_{j \geq j'} \sum_{s \in \mathbb{Z}} d_{j,s} \psi_{j,s}(x) \quad , \quad (5)$$

where the coefficients are obtained from the inner products

$$c_{j',s} = \int_{-\infty}^{+\infty} f(x) \phi_{j',s}(x) dx, \quad d_{j,s} = \int_{-\infty}^{+\infty} f(x) \psi_{j,s}(x) dx. \quad (6)$$

The coefficients $c_{j',s}$ represent a “smooth” approximation of $f(x)$ up to the scale $2^{-j'}$ (where, in general, $j'=0$ so that it corresponds to the largest scale) and the second sum in (5) adds up details at all smaller scales. Implicitly stated by equation (5) is the fact that the complete wavelet basis consists of the space $V_{j'}$ spanned by discrete translations of the averaging function at scale j' and the sum of all subspaces W_j which are spanned by translations of the wavelet functions at the respective scales $j \geq j'$.

2.2 Periodic wavelet basis

For use with periodic functions with a period of 1, the bases of the MRA take the following form (Perrier & Basdevant, 1989):

$$\begin{aligned} \phi_s^j(x) &= \sum_{z \in \mathbb{Z}} \phi_{j,s}(x+z) \\ \psi_s^j(x) &= \sum_{z \in \mathbb{Z}} \psi_{j,s}(x+z) \end{aligned} \quad j \geq 0, \quad 0 \leq s \leq 2^j. \quad (7)$$

The decomposition becomes:

$$f(x) = c_0^0 + \sum_{j \geq 0} \sum_{0 \leq s < 2^j} d_s^j \psi_s^j(x) \quad , \quad (8)$$

with

$$c_0^0 = \int_0^1 f(x) dx, \quad d_s^j = \int_0^1 f(x) \psi_s^j(x) dx. \quad (9)$$

An important property is the conservation of energy during the transform, i.e.

$$\int_0^1 |f(x)|^2 dx = |c_0^0|^2 + \sum_{j \geq 0} \sum_{s=0}^{2^j-1} |d_s^j|^2 \quad . \quad (10)$$

2.3 Multi-dimensional wavelet basis

In order to perform a multi-dimensional MRA, we will use tensor products of one-dimensional wavelet bases. In $n = 3$ (2) dimensions there need to be $Q(n) = 7$ (3) orthogonal subspaces $W_{q,j}$ per scale which together with the subspace V_j make up a complete basis (Mallat, 1989; Meneveau, 1991). Actually, the number Q can be understood as the number of coordinate directions and possible diagonals that can be defined in the n -dimensional space (i.e. in three dimensions there are three coordinates, one diagonal and its three projections upon the different planes). The n -dimensional decomposition reads:

$$f(\mathbf{x}) = c_{0,\dots}^0 + \sum_{j \geq 0} \sum_{q=1}^{Q(n)} \sum_{s,\dots=0}^{2^j-1} d_{s,\dots}^{q,j} \psi_{s,\dots}^{q,j}(\mathbf{x}) \quad , \quad (11)$$

where

$$d_{s,\dots}^{q,j} = \int_0^1 f(\mathbf{x}) \psi_{s,\dots}^{q,j}(\mathbf{x}) d\mathbf{x}. \quad (12)$$

As an example, the basic averaging and wavelet functions are below given in the two-dimensional case:

$$\phi_{s,l}^j(x, y) = \phi_s^j(x) \cdot \phi_l^j(y) \quad \rightarrow \quad \text{subspace } V_j \quad (13a)$$

$$\psi_{s,l}^{1,j}(x, y) = \phi_s^j(x) \cdot \psi_l^j(y) \quad \rightarrow \quad W_{1,j} \quad (13b)$$

$$\psi_{s,l}^{2,j}(x, y) = \psi_s^j(x) \cdot \phi_l^j(y) \quad \rightarrow \quad W_{2,j} \quad (13c)$$

$$\psi_{s,l}^{3,j}(x, y) = \psi_s^j(x) \cdot \psi_l^j(y) \quad \rightarrow \quad W_{3,j}, \quad (13d)$$

where $\psi_{s,l}^{1,j}$ selects details in the horizontal direction, $\psi_{s,l}^{2,j}$ along the vertical and $\psi_{s,l}^{3,j}$ along the diagonal.

2.4 Interpolation step

In practice, the signal to be analyzed will be available in discrete form rather than analytically. We suppose a regular grid of dimension 2^{nJ} where functional values $f[\mathbf{i}]$ are

known (\mathbf{i} being the n -dimensional vector of indices to the grid nodes). The very first thing to do in our MRA is to determine an approximation \tilde{f} which belongs to the space V_J , i.e. the space in which work is done at the finest level of the hierarchy. Following Perrier & Basdevant (1989), \tilde{f} is obtained by filtering the data $f[\mathbf{i}]$ with a function L_J given by

$$\hat{L}_J(k) = \frac{1}{2^J} \frac{\hat{S}(k)}{\hat{\phi}_0^J(k)} \quad 0 \leq k < 2^J, \quad (14)$$

where S is the cardinal Lagrange function which verifies

$$S\left(\frac{m}{2^J}\right) = \delta_{0,m} \quad m = 0 \dots 2^J - 1, \quad (15)$$

and the Fourier coefficients are defined as

$$\hat{g}(k) = \int_0^1 g(x) e^{-2i\pi kx} dx, \quad k \in \mathbb{Z}, \quad (16)$$

with the restitution formula

$$g(x) = \sum_{k \in \mathbb{Z}} \hat{g}(k) e^{2i\pi kx}. \quad (17)$$

As remarked by Do-Khac *et al.* (1994), the filter L_J is nearly a Dirac filter and often not applied in practice, which, however, can lead to misinterpretations of the smallest scales of the transform. More specifically, the interpolation by L_J assures that the final decomposition (e.g. equ. (11)) will be an interpolation function with respect to the discrete values $f[\mathbf{i}]$. In the present work, we do perform the interpolation step consistently.

2.5 Mallat's fast algorithm: Quadrature mirror filters

The fast transform between physical and wavelet space relies upon an efficient evaluation of the required scalar products. Coefficients at a given scale are recursively determined by convolution of previous coefficients with either one of two quadrature mirror filters (QMF) H or G and subsequent down-sampling by a factor of two, e.g. in the one-dimensional case:

$$\begin{aligned} c_k^{j-1} &= \sum_{s=0}^{2^j-1} c_s^j \cdot H_j(s-2k) \\ d_k^{j-1} &= \sum_{s=0}^{2^j-1} d_s^j \cdot G_j(s-2k) \end{aligned} \quad 0 \leq k < 2^{j-1}. \quad (18)$$

Due to the properties of the MRA, the filters are directly linked to the basic averaging and wavelet functions. In the periodic case, these relations read (Perrier & Basdevant, 1989):

$$\hat{H}_j(k) = \frac{\sqrt{2}}{2^j} \frac{\hat{\phi}(2k/2^j)}{\hat{\phi}(k/2^j)}, \quad \hat{G}_j(k) = \frac{\sqrt{2}}{2^j} \frac{\hat{\psi}(2k/2^j)}{\hat{\phi}(k/2^j)}. \quad (19)$$

The sequence of operations, consisting of convolutions and down-sampling, is schematically illustrated in figure 1. At each level j , an array of $2^{n(j+1)}$ coefficients c^{j+1} is

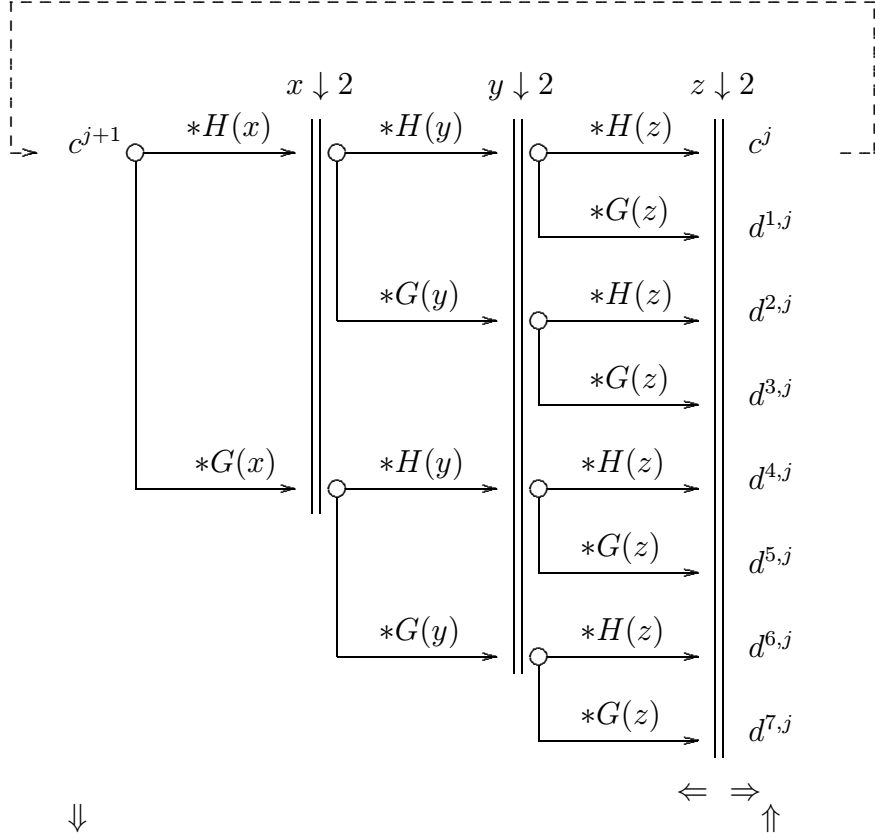


Figure 1: Schematic of the j th step of Mallat's fast algorithm for wavelet decomposition of an $n=3$ dimensional data field; '*' represents the convolution operator, ' $x_i \downarrow 2$ ' signifies down-sampling by a factor of two in the direction x_i . c^j is an array of dimension 2^{3j} corresponding to the "smooth" content of the data in space V_j and the equally dimensioned arrays $d^{q,j}$, where $1 \leq q \leq 7$, hold the "detail" information of the signal at the current scale j and respective direction q . Supplementary symbols in the lower part of the figure indicate the occurrence of operations which pertain to the execution on distributed memory machines (cf. §3.1 and 3.3): \leftarrow , transpose from 'z-cut' to 'y-cut' representation of data; \Rightarrow , inverse transpose; \uparrow , up-sizing the number of active processors; \downarrow , down-sizing.

processed, thereby producing an array of 2^{nj} new coefficients c^j (for further use in the following step) and a number of Q arrays of wavelet coefficients $d^{q,j}$ each of the same size 2^{nj} . Therefore, the full transform of 2^{nJ} discrete data values leads to $1 + Q \sum_{j=0}^{J-1} 2^{nj} = 2^{nJ}$ wavelet coefficients.

The inverse transform (“wavelet synthesis”) proceeds in opposite direction along the same schematic with the following modifications (e.g. Farge, 1992, p. 418): (i) replacing convolutions with de-convolutions; (ii) performing up-sampling instead of down-sampling; (iii) wherever two “paths” join in figure 1, the signal is added and multiplied by a factor of 2.

2.6 Processing in Fourier space: Non-compact filters

Selecting a type of wavelet means, for instance, to chose an appropriate filter $H(x)$ and to determine the corresponding complement filter via $G(l) = (-1)^{1-l}H(1-l)$. The choice is based upon criteria concerning fundamental properties of wavelets (Farge, 1992):

- (i) Regularity. The analyzing wavelet should be localized (i.e. should have low intensity outside a “small” region) in physical and Fourier space. By this virtue, a signal can be analyzed in space and scale simultaneously.
- (ii) Cancellations. In order for the analysis to pick up high-order fluctuations of the signal reliably, the analyzing wavelet should possess a corresponding number of vanishing high-order moments.

Cubic spline wavelets have been widely used for the analysis of turbulent flow fields (e.g. Meneveau, 1991; Do-Khac *et al.*, 1994) offering a reasonable compromise with respect to localization and smoothness (4 vanishing moments). Since the associated filters are not compactly supported, it is in this case advantageous to work in Fourier space using FFT’s. For the present study, we have chosen such cubic spline functions as our analyzing functions. The expressions for computing the filters H , G , L are given in appendix A. The reader is referred to Fröhlich & Schneider (1994) who give a detailed account of the decomposition algorithm including the down/up-sampling step in Fourier space (cf. their §5.1).

3 Parallel algorithm for three-dimensional discrete wavelet decomposition

When dealing with large, three-dimensional data sets, the operation count for the discrete wavelet transform (DWT) – which scales as $\mathcal{O}(N^3 \log(N))$ – can become unacceptable. An even more serious problem is the required core memory which easily exceeds current possibilities. Therefore, we distribute the work load across various processors of a distributed memory machine and handle communication via MPI.

Different algorithms for parallelizing the DWT have been proposed in the literature (e.g. Møller Nielsen, 1998; Yang & Misra, 1998)¹. Those authors are concerned with

¹Our literature research is admittedly far from being exhaustive.

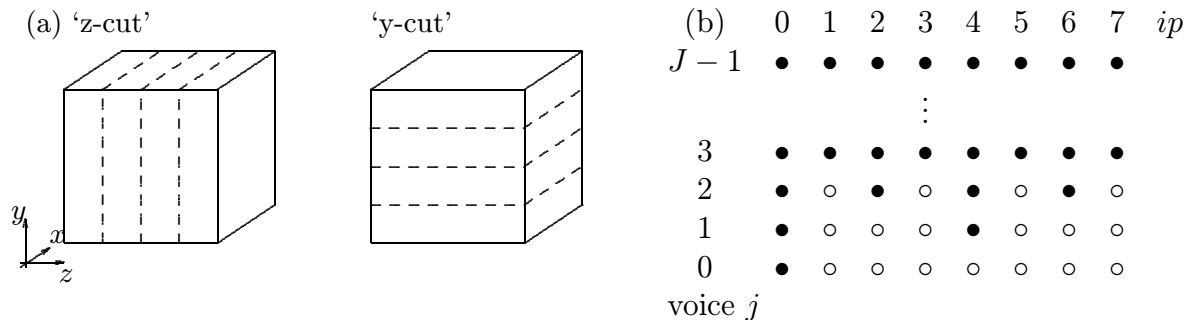


Figure 2: (a) The ‘slice’ data-model consisting of complete planes resident in local memory and therefore requiring one pair of transpose operations for performing ‘global’ operations (such as fast Fourier transforms) in all three directions. (b) The hierarchy of ‘active’ processors (indicated by full symbols) during the coarsest levels of Mallat’s algorithm. The scheme shows an example with 8 available processors; $N=2^J$ is the problem size and ip is the global identification of each processor.

compact filters which lead to very different algorithms since almost local operations are performed on the data. Compactly supported wavelets, however, are in general not symmetric which reduces their usefulness for flow analysis. This is one of the reasons why we chose non-compact spline wavelets instead.

The algorithm exposed in the following is especially adapted for the use of long filters and “global” operations like fast Fourier transforms. It uses a “slice” data model as is commonly used in spectral codes for fluid mechanics. We should stress that our main concern in designing the parallelization was memory and not necessarily execution speed because we intend to use wavelets as tools for analysis and not for the construction of simulation methods.

In order to simplify data dependencies during the dyadic algorithm, we make the important assumption that the number of active processors is a power of two: $n_P = 2^P$.

3.1 The ‘slice’ model for data in physical or Fourier space

Three-dimensional data arrays will be stored across different processors by partitioning either along the z -coordinate or along the y -coordinate (figure 2 a) in different parts of the code. Complete (x, y) - or (x, z) -planes are therefore resident in memory which allows to perform full vector operations in both those directions without data exchange (cf. table 1 for the dimensions of distributed arrays in physical and Fourier space).

As we have seen in §2, the three elementary operations of the transform are:

- (i) FFT;
- (ii) down/up-sampling;
- (iii) convolution.

type	global dimension	local dimension	
		z-cut	y-cut
real	$(n_x + 2, n_y, n_z)$	$(n_x + 2, n_y, \frac{n_z}{n_p})$	$(n_x + 2, \frac{n_y}{n_p}, n_z)$
complex	$(0 : n_x/2, 0 : n_y - 1, 0 : n_z - 1)$	$(0 : n_x/2, 0 : n_y - 1, \frac{n_z}{n_p})$	$(0 : n_x/2, \frac{n_y}{n_p}, 0 : n_z - 1)$

Table 1: The dimensions of three-dimensional data arrays distributed over a number of n_p processors in either ‘z-cut’ or ‘y-cut’ representation. Note that the actual data has dimension $\text{real}(n_x n_y n_z)$ in physical space and additional “odd-ball” indices are due to the real-to-complex Fourier transform. Also, it should be mentioned that the dimensions should be considered variable since, due to down/up-sampling, the lengths vary continuously during the algorithm, i.e. $1 \leq n_x, n_y, n_z \leq 2^{J-1}$. Therefore, data sizes and distribution across processors (which also vary in number, $n_p \neq n_P$ in general) must be recalculated at various stages.

All three operations (i)-(iii) are factorized into individual spatial directions. Points (i) and (ii) require non-local data-access. Therefore, the data base has to be transposed and back-transposed once per step j in order to switch between ‘z-cut’ and ‘y-cut’ representation. The occurrence of the transpose is indicated by horizontal arrows in the schematic of the algorithm shown in figure 1.

The parallel algorithm adds basically two features to the above list of operations:

- (iv) transpose of distributed data;
- (v) resize: decrease/increase the number of active processors.

The technique for performing (iv) will be explained in §3.2 below. The feature (v) is necessary because with any initial number of active processors $n_P > 1$, and supposing that the decomposition is carried out for all voices $j = 0 \dots J - 1$, the algorithm will reach a point where the data size in one direction is smaller than n_P . In those instances, a redistribution of data and a de-activation of a subset of processors will be carried out before continuing the computation. This point will be discussed in §3.3. Finally, in §3.4 we will present the way wavelet coefficients are stored in distributed fashion and how they can be accessed.

3.2 Transpose of the distributed data base

Transposing a distributed array A , dimensioned as in table 1, from ‘y-cut’ to ‘z-cut’ representation is carried out via the following scheme:

```

for  $ip = 0$  to  $n_p$  do
  if  $ip$  is active then
    if  $ip \neq myid$  then
      • send  $A(:, jbeg(myid) : jend(myid), kbeg(ip) : kend(ip))$  to processor  $ip$ 
      • receive from  $ip$  into  $work(:)$ 

```

```

else
  work(:) ← A(:, jbeg(myid) : jend(myid), kbeg(myid) : kend(myid)) {locally}
end if
A(:, jbeg(ip) : jend(ip), kbeg(myid) : kend(myid)) ← work
end if
end for

```

where $0 \leq myid \leq n_P - 1$ is the identification attributed to a processor; $jbeg$, $jend$, $kbeg$, $kend$ are the pointers indicating the range of *global* indices which are held by each processor in ‘y-cut’ and ‘z-cut’ representation respectively. The inverse transpose operation proceeds in analog fashion.

Three-dimensional Fourier transforms are performed separately along each spatial direction at the additional cost of a pair of transpose operations before and after treating the *z*-direction.

3.3 Resize: Changing the number of active processors

Whenever the data length of the current scale drops below the number of currently active processors, i.e. $2^{j-1} < n_p$, the latter needs to be reduced. One possibility – and probably the simplest algorithmically – is to switch directly to single-processor mode at this point. However, it means to leave $(n_P - 1)/n_P$ of the resources idle for the remainder of the decomposition. With a reasonable amount of extra coding effort, half of the active processors can be carried along at each successive step.

Since the algorithm is dyadic, n_p will also be resized by a factor of two each time until – in the last step – only one processor remains active. This sequence, where at each down-sizing step one active processor out of two is de-activated, is depicted in figure 2 (b). The inverse operation of up-sizing during the inverse wavelet transform proceeds in reverse order along the same diagram.

The important part of resizing is the data exchange between ‘survivors’ (‘veterans’) and newly de-activated (re-activated) processors. During down-sizing, de-activating processors means making them send the coefficients c^j – needed for carrying on with the decomposition – to their ‘surviving’, lower-indexed neighbour. Analogously, processors which are re-activated during up-sizing receive half of the coefficients c^j from their inferior ‘veteran’ neighbours.

Resizing is always carried out when data is arranged in ‘z-cut’ representation (cf. figure 1 for the occurrence during the algorithm) which allows exchanged data blocks to be sent/received in-place and no intermediate storage is necessary. The following pseudo-code shows how this communication can be handled:

```

if receiving(myid) = true then
  if down-sizing then
    n ← (n_x/2 + 1)n_y(kend(myid) – kbeg(myid) + 1)
    nl ← n + 1
  else if up-sizing then
    n ← (n_x/2 + 1)n_y(kend(irecv) – kbeg(irecv) + 1)
  end if
end if

```

```

     $nl \leftarrow 1$ 
end if
    • receive  $n$  coefficients from irecv starting with local address  $nl$ 
else if  $sending(myid) = true$  then
    if down-sizing then
         $n \leftarrow (n_x/2 + 1)n_y(kend(myid) - kbeg(myid) + 1)$ 
         $nl \leftarrow 1$ 
    else if up-sizing then
         $n \leftarrow (n_x/2 + 1)n_y(kend(myid) - kbeg(myid) + 1)/2$ 
         $nl \leftarrow n + 1$ 
    end if
    • send  $n$  coefficients to isend starting with local address  $nl$ 
end if

```

Here, the pointers *irecv* and *isend* refer to indices of the pair of processors which is mutually exchanging data and has been set according to the tree-structure indicated in figure 2 (b). Note that the indices and buffer sizes correspond to data in Fourier space, i.e. of complex type (cf. table 1).

3.4 Distributed storage of wavelet coefficients

Due to the fact that different processors participate in the computation of a different number of levels of the dyadic decomposition, the number of wavelet coefficients $d_{s,l,k}^{q,j}$ generated by each one of them varies. This memory unbalancing was of some concern when laying out the present method and it was initially planned to equilibrate the use of resources by additional data exchange at the end of each step. As we will see in the following, however, the actual difference is quite small and we content ourselves with a slight over-dimensioning of the arrays.

The overall data-volume produced by a single processor is given by the following formula:

$$N_{ip} = 7 \cdot \underbrace{\sum_{j=j_1}^{J-1} \frac{2^{3j}}{n_P}}_{N_1} + 7 \cdot \underbrace{\sum_{j=j_2}^{j_1-1} 2^{2j}}_{N_2(j_2)} \quad , \quad (20)$$

where N_1 is the data-volume accumulated by all processors before the first down-sizing which occurs at $j = j_1 = P$ and N_2 is the remainder up to the last voice j_2 at which the respective processor remains active. This value j_2 varies from processor to processor in the range $0 \leq j_2 \leq j_1$. The maximum difference in data volume is given by $N_2(j_2=0)+1$ which only depends on the number of processors n_P and not on the actual problem size:

$$N_2(j_2=0) = (n_P^2 - 1) \cdot \frac{7}{3} \quad . \quad (21)$$

Accordingly, the data unbalancing in a case of $n_P = 128$ processors consists of only 5462 real elements (corresponding to 21848 Byte in a 32 bit representation). This disadvantage seems acceptable. We therefore need to allocate the following size for the buffers which

accommodate the transform:

$$N_{DWT} = N_1 + N_2(j_2=0) + 1 = \frac{2^{3J}}{n_P} + \frac{4n_P^2 - 7}{3} + 1 \quad . \quad (22)$$

In the present scheme, coefficients are stored sequentially, voice by voice, and looping over the directional index ‘ q ’. For a given set of indices $\{s, l, k, q, j\}$, the processor identification and the local index of the corresponding coefficient $d_{s,l,k}^{q,j}$ can be determined by the following procedure:

Require: s, l, k, q, j

$n_j \leftarrow 2^j$

$n_t \leftarrow \max(0, P - j)$ {the number of resizes so far}

$n_p \leftarrow 2^{P-n_t}$ {the current number of active processors}

for $ip = 0$ to $n_P - 1$ **do**

if $\text{mod}(ip, n_P/n_p) = 0$ **then**

$isactive(ip) \leftarrow 1$

else

$isactive(ip) \leftarrow 0$

end if

end for

if $isactive(myid) = 1$ **then**

for $i = 0$ to $myid$ **do**

$n_{inf} \leftarrow n_{inf} + isactive(i)$ {number of active procs with inferior address}

end for

$k_f \leftarrow (n_{inf} - 1)(n_j/n_p) + 1$ {first z -index held locally}

$k_l \leftarrow n_{inf}(n_j/n_p)$ {last z -index held locally}

if $k \geq k_f$ and $k \leq k_l$ **then**

if $n_t > 0$ **then**

$n_{sum} \leftarrow 2^{3J}/n_P - n_P^2 + 7(n_P^2 - 2^{j+1})/3$

else

$n_{sum} \leftarrow (2^{3J}/n_P - 2^{3(j+1)})/n_P$

end if

$index \leftarrow n_{sum} + (q - 1)(k_l - k_f + 1)n_jn_j + (k - k_f)n_jn_j + (l - 1)n_j + s$

\Rightarrow processor $myid$ holds coefficient $d_{s,l,k}^{q,j}$ at local position $index$

 • broadcast the value of $myid$ and $index$ to global variables $ipholds$ and $index$

end if

end if

The retrieval of the actual value of the coefficient will then be performed by a request to processor $ipholds$ at address $index$.

4 Some results

In this section we present some results from wavelet transforms of two-dimensional and three-dimensional data as well as measures of the performance of our parallel algorithm.

4.1 Two-dimensional DWT

Figure 3 shows the transform of an image showing a fingerprint from Bradley, Brislawn & Hopper (1993)². The wavelet coefficients are given in the usual scale-wise and direction-wise representation, each block $d^{q,j}$ being separated by a solid line. By comparing the coefficients with the original signal, it can be seen that the $d^{1,j}$'s (upper left block) pick up variations in the horizontal direction, $d^{2,j}$'s (lower right block) react to vertical variations and $d^{3,j}$'s (upper right block) to diagonal variations. The data has then been truncated in wavelet space, i.e. all coefficients with an intensity below a certain threshold have been set to zero, and back-transformed to physical space. As is shown in the same figure, the reconstruction with only 1.5% of the total number of coefficients gives a very reasonable image (containing 98% of the “energy”). The trace map of the intense coefficients shows that all scales contribute but in a spatially inhomogeneous way. This is one of the reasons why wavelet compression is often more efficient than compression based upon global approximations, e.g. cosine transforms.

Figure 4 gives a similar impression in the case of turbulent flow data. The two-dimensional transform has been applied to instantaneous enstrophy values of a slice through statistically homogeneous-isotropic flow obtained by direct numerical simulation (DNS, cf. §4.2). It is evident – and representative of many snapshots – that at small scales, high intensity coefficients appear in isolated regions, hinting at the intermittent nature of turbulence.

4.2 Three-dimensional DWT

The data used in this section is taken from a large-scale forced simulation of homogeneous-isotropic, turbulent flow at $Re_\lambda = 45$ using a pseudo-spectral code with 128^3 de-aliased Fourier modes.

First, we have tested the compression capabilities of the spline base. Farge, Schneider & Kevlahan (1998) have proposed a wavelet based technique for separating the flow field into a coherent part and an incoherent background flow. The procedure is simply a truncation of vorticity in wavelet space according to an “objective” threshold C_T which only depends upon the variance of the signal and the number of samples, viz.

$$c_T = (2 Z \log_{10} (N^n))^{1/2} \quad , \quad (23)$$

where the total enstrophy is defined as

$$Z = \frac{1}{2} \int_{\Omega} |\boldsymbol{\omega}|^2 dx \quad . \quad (24)$$

Applying this criterion, approximately 99.8% of the wavelet coefficients are discarded in the instantaneous flow fields of our simulation. Figure 5 shows an example of the reconstruction, where 86% of enstrophy is carried by the remaining few coefficients. The

²(also contained in the MATLAB package WaveLab 802, available from www-stat.stanford.edu)

figure shows that most of the regions of strong vorticity, i.e. the “worms”, are preserved and can be identified by comparison with the full field.³

The energy spectral tensor has been defined by Meneveau (1991) as follows:

$$E_{ij}(k_m) = \frac{2^{2(J-(j-1))} h}{2\pi \log(2)} \sum_{q=1}^7 \sum_{s,k,l=1}^{2^m} d_{s,k,l}^{q,m}(u_i) \cdot d_{s,k,l}^{q,m}(u_j) \frac{1}{2 N^3} \quad , \quad (25)$$

where h is the discrete grid size and $d_{s,k,l}^{q,m}(u_i)$ are the wavelet coefficients of the transform of the velocity component u_i . The three-dimensional energy spectrum $E(k_m)$ is then simply the trace of (25). Figure 6 shows the wavelet spectrum for the flow field used before in figure 5. It is obvious that the spacing of the wavenumbers is very coarse. Furthermore, deviations from the usual radial Fourier spectrum are visible, due to the non-locality of the basis functions in Fourier space. The use of the wavelet spectrum – when averaged over the whole domain – is therefore limited. However, the fact that local spectra can be studied is potentially useful for the analysis of inhomogeneous fields (Do-Khac *et al.*, 1994).

4.3 Performance of the parallel DWT

For realistic problem sizes, the memory requirements of the DWT decrease near linearly with the number of processors, cf. (22). With this property, the algorithm fulfills our most important initial specification. On the other hand, it is interesting to have a look at the execution time using different numbers of processors.

We have used the following code structure for timing the transform:

```

c      /* initialize MPI */
c      call MPI_INIT(ierr)
c      call MPI_COMM_RANK(MPI_COMM_WORLD,myid,ierr)
c      call MPI_COMM_SIZE(MPI_COMM_WORLD,numprocs,ierr)
c      /* start timing */
c      t0=MPI_WTIME()
c      /* do forward wavelet transform */
c      call dwt3dp(signal,+1)
c      t1=MPI_WTIME()
c      /* backward DWT */
c      call dwt3dp(signal,-1)
c      t2=MPI_WTIME()
c      dt1=(t1-t0)
c      dt2=(t2-t1)
c      call MPI_REDUCE(dt1,dt1m,1,MPI_DOUBLE_PRECISION,
$      MPI_MAX,0,MPI_COMM_WORLD,ierr)
c      call MPI_REDUCE(dt2,dt2m,1,MPI_DOUBLE_PRECISION,
```

³A short animation of a sequence of 50 such fields at a temporal resolution of one Kolmogorov time unit can be found under www.pik-potsdam.de/~uhlmann/reports/homiso/compress.html in MPEG format.


```

$      MPI_MAX,0,MPI_COMM_WORLD,ierr)
      if(myid.eq.0)write(*,*)'max: ',dt1m+dt2m
c      /* finalize mpi */
      call MPI_FINALIZE(ierr)

```

The execution time is therefore the total elapsed time while performing a forward and a backward wavelet transform.

Our tests have been carried out on an IBM SP2 with nodes at 66 MHz clock speed and 128 MB core memory as well as a CRAY T3E LC 384 with DEC Alpha EV5.6 nodes at 600MHz clock speed and 512 MB core memory. Figures 7 and 8 show the speed-up and the parallel efficiency obtained when working on a field with dimension $N = 128$ and $N = 256$, respectively. The algorithm is scaling reasonably well.

5 Conclusions

We have employed n -dimensional periodic wavelet bases as suggested by (Perrier & Basdevant, 1989; Meneveau, 1991; Do-Khac *et al.*, 1994). We have devised and implemented a fast algorithm on distributed memory machines using a ‘slice’ data model and restricting to a number of processors which is a power of two. This latter restriction is not strictly mandatory but simplifies data dependencies. It would be straightforward to construct the algorithm such that the transform works with an arbitrary number n_P until data becomes scarce for the first time (i.e. until $2^{j-1} < n_P$) and then proceed with the nearest inferior power of two (e.g. $n_P = 7$, then continue with $n_p = 4$).

The present algorithm satisfies our intention of reducing local memory requirements near linearly. Furthermore, execution time was seen to speed up very reasonably.

Source code written in FORTRAN and using MPI can be obtained from the author upon request. Computing time (on the CRAY T3E system) provided by the computer center at ZIB, Berlin, is gratefully acknowledged.

References

- BRADLEY, J., BRISLAWN, C. & HOPPER, T. 1993 The FBI wavelet/scalar quantization standard for gray-scale fingerprint image compression. In *Proc. SPIE*, , vol. 1961.
- BRASSEUR, J. & WANG, Q. 1992 Structural evolution of intermittency and anisotropy at different scales analyzed using three-dimensional wavelet transforms. *Phys. Fluids A* **4** (11), 2538–2554.
- DAUBECHIES, I. 1992 *Ten lectures on wavelets*. Philadelphia, PA, USA: SIAM.
- DO-KHAC, M., BASDEVANT, C., PERRIER, V. & DANG-TRAN, K. 1994 Wavelet analysis of 2d turbulent fields. *Physica D* **76**, 252–277.
- FARGE, M. 1992 Wavelet transforms and their applications to turbulence. *Ann. Rev. Fluid Mech.* **24**, 395–457.

- FARGE, M., SCHNEIDER, K. & KEVLAHAN, N. 1998 Coherent structure eduction in wavelet-forced two-dimensional turbulent flows. In *IUTAM Symposium on Dynamics of Slender Vortices* (ed. E. Krause & K. Gersten), pp. 65–83.
- FROEHLICH, J. 2000 LES of vortex shedding past circular cylinders. In *ECCOMAS*. Barcelona, Spain.
- FRÖHLICH, J. & SCHNEIDER, K. 1994 An adaptive wavelet Galerkin algorithm for one- and two-dimensional flame computations. *Eur. J. Mech. B/Fluids* **13** (4), 439–471.
- KLEIN, R. 1998 Kleinskalige Instabilitäten als Bausteine der turbulenten Energiekaskade. DFG-Antrag auf Gewährung einer Sachbeihilfe, Projekt KL 611/10 (in German).
- MALLAT, S. 1989 A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Analysis Mach. Intell.* **11** (7), 674–693.
- MENEVEAU, C. 1991 Analysis of turbulence in the orthonormal wavelet representation. *J. Fluid Mech.* **232**, 469–520.
- MEYER, Y. 1992 *Wavelets and operators*. Cambridge Univ. Press.
- MØLLER NIELSEN, O. 1998 Parallel wavelet transforms. In *Applied Parallel Computing. Large Scale Scientific and Industrial Problems* (ed. J. W. et al.), *Lecture Notes in Computer Science*, vol. 1541, pp. 385–389. Springer.
- PERRIER, V. & BASDEVANT, C. 1989 Periodical wavelet analysis, a tool for inhomogeneous field investigation. Theory and algorithms. *Rech. Aérop.* **3**, 54–67.
- TORRENCE, C. & COMPO, G. 1998 A practical guide to wavelet analysis. *Bull. Am. Met. Soc.* **79** (1), 61–78.
- YANG, L. & MISRA, M. 1998 Coarse-grained parallel algorithms for multi-dimensional wavelet transforms. *J. Supercomputing* **12**, 99–118.

A Filter functions associated with spline wavelets

It is convenient to assign the filters directly in Fourier space. Using spline functions of order m , the expressions read (Perrier & Basdevant, 1989):

$$\hat{H}_j(k) = \frac{1}{2^j} H\left(\frac{k\pi}{2^j}\right), \quad \text{where} \quad (26a)$$

$$H(x) = \sqrt{2} \cos^m(x) \left(\frac{P_{m-1}(\sin^2(x))}{P_{m-1}(\sin^2(2x))} \right)^{1/2}, \quad (26b)$$

$$\hat{G}_j(k) = \frac{1}{2^j} G\left(\frac{k\pi}{2^j}\right), \quad \text{where} \quad (27a)$$

$$G(x) = \sqrt{2} e^{-2ix} \sin^m(x) \left(\frac{P_{m-1}(\cos^2(x))}{P_{m-1}(\sin^2(2x))} \right)^{1/2}, \quad (27b)$$

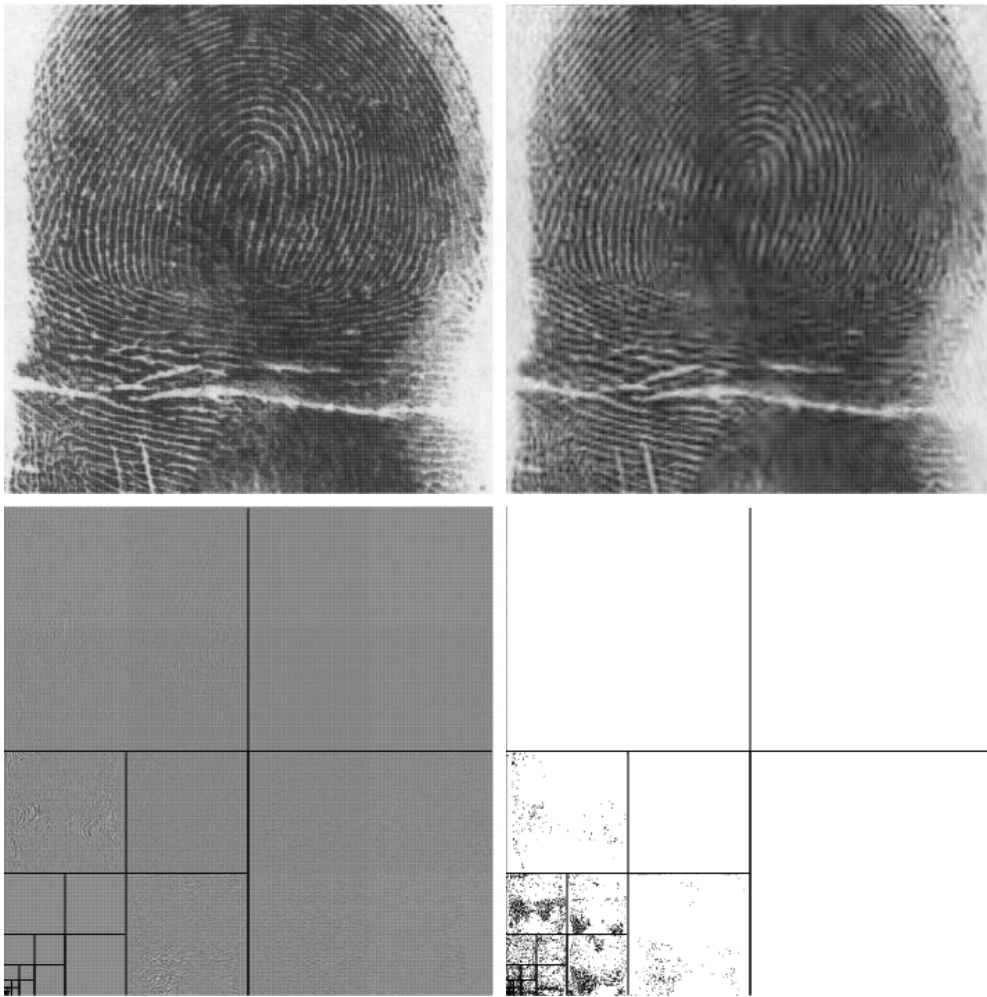


Figure 3: Wavelet transform of a 512×512 , 8 bit per pixel, gray-scale image showing a fingerprint (taken from Bradley *et al.* 1993). In the upper left, the original data; upper right, its reconstruction after keeping only the most intense 1.5% of the coefficients. The lower left shows the wavelet coefficients, where small-scales are over-exposed for clarity. In the lower right, the scale-wise and direction-wise distribution of the strongest 1.5% of the coefficients kept in the reconstruction is shown. Scales j and directions q are separated by solid lines. Note that these images (and the following) are reproduced on paper in low resolution for practical purposes; high-resolution images are available online under the URL www.pik-potsdam.de/~uhlmann/reports/homiso/wavelets.html.

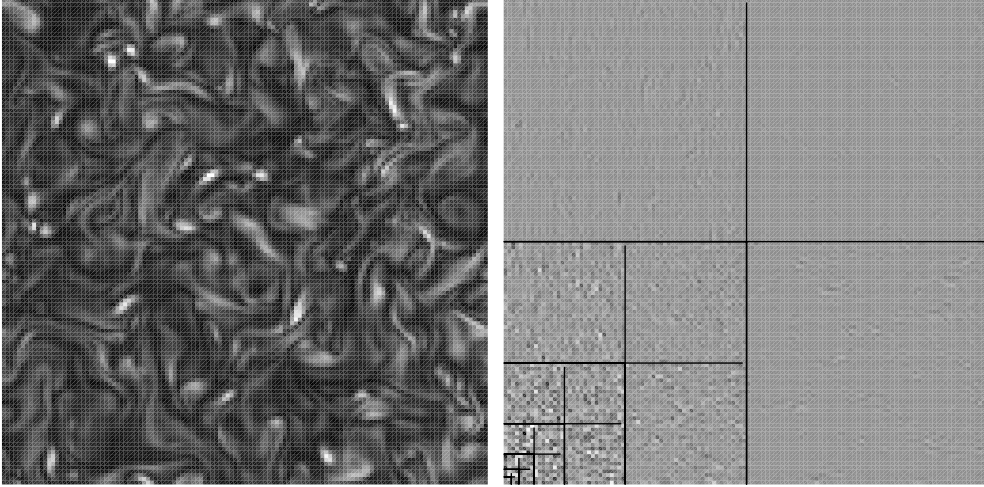


Figure 4: Wavelet transform of a slice through a forced, 128^3 modes DNS of homogeneous isotropic turbulence at $Re_\lambda = 45$. On the left, the original data. The image on the right shows the intensity of the wavelet coefficients.

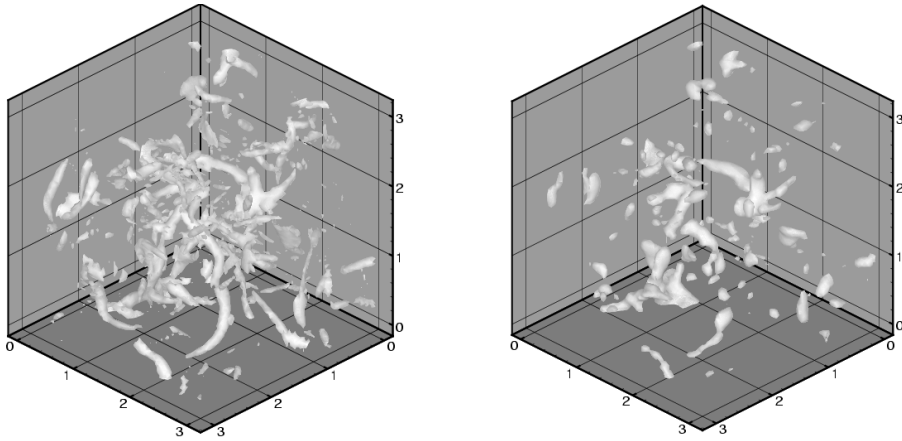


Figure 5: Isosurfaces of vorticity magnitude at $|\boldsymbol{\omega}| = 4\omega'$ of homogeneous-isotropic flow at $Re_\lambda = 45$. The linear dimension of the visualized sub-domain corresponds to 225 Kolmogorov lengths. On the left, the full field; on the right, the field which is reconstructed after high-pass filtering in wavelet space according to the threshold given in (23) where 99.2% of the coefficients are discarded.

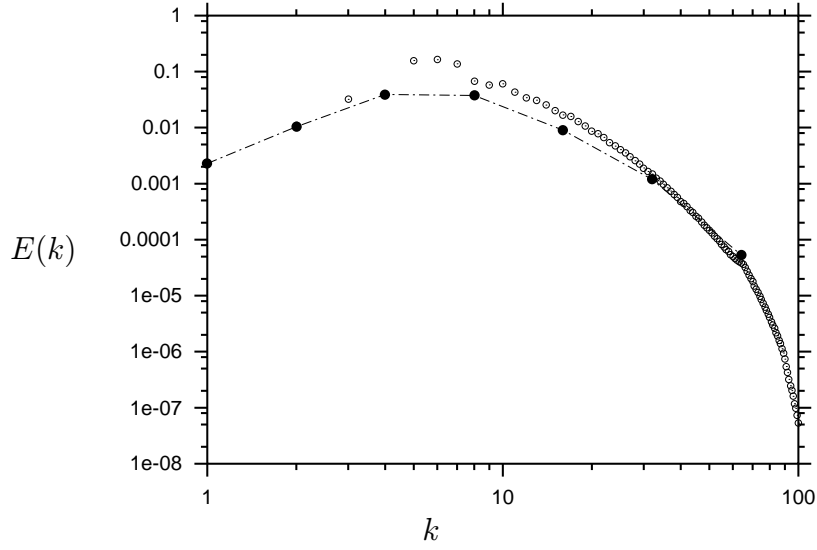


Figure 6: Three-dimensional energy spectrum of a snapshot of large-scale forced, homogeneous-isotropic flow at $Re_\lambda = 45$. \bullet , wavelet spectrum; \circ , Fourier spectrum.

$$\hat{L}_J(k) = \frac{1}{2^{3J/2}} L\left(\frac{k\pi}{2J}\right), \quad \text{where} \quad (28a)$$

$$L(x) = \frac{(P_{m-1}(\cos^2(x)))^{1/2}}{P_{m/2-1}(\sin^2(x))}. \quad (28b)$$

The coefficients of the spline polynomials $P_m(x) = \sum_{i=0}^m a_i^m x^i$ can be obtained by the following recursive formulas:

$$a_0^1 = -2/3, \quad a_0^j = 1 \quad j = 1 \dots m, \quad (29)$$

$$a_i^j = \frac{1}{j(2j+1)} \left((j-1)(2j-2i+1) a_i^{j-1} - 2(j-i+1)^2 a_{i-1}^{j-1} \right) \quad 1 \leq i \leq m.$$

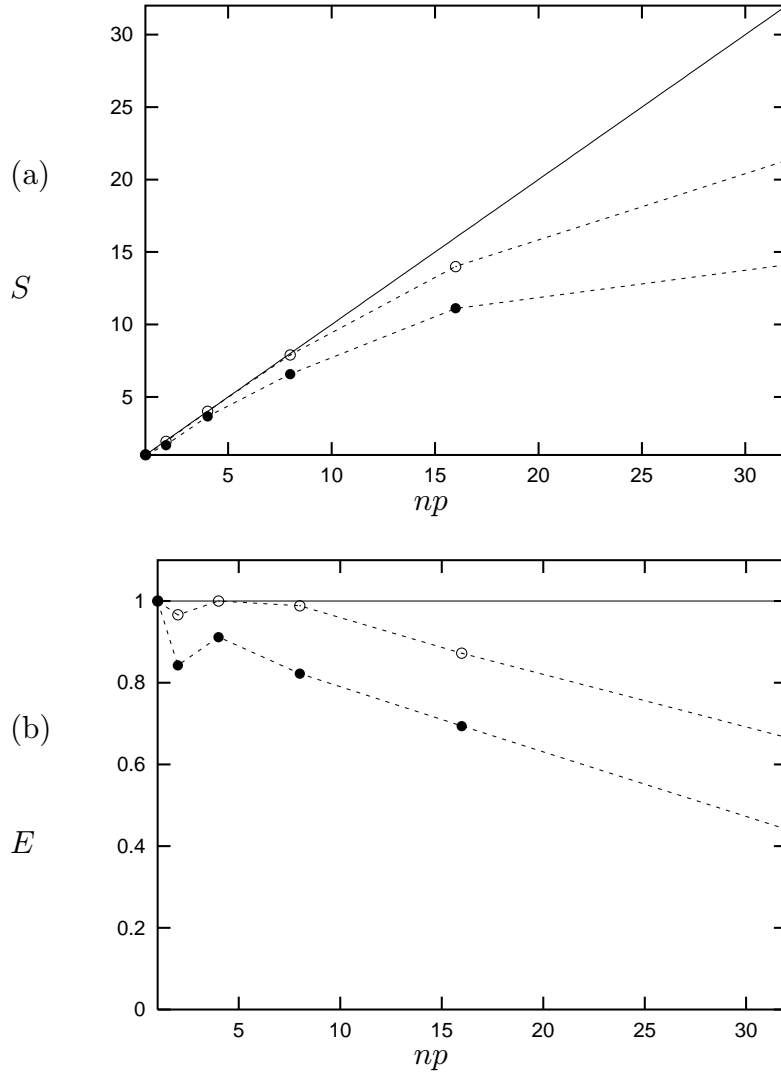


Figure 7: Cumulated execution times of wavelet transform and inverse transform of a field of dimension 128^3 : (a) relative speed-up $S(np) = \Delta t(1)/\Delta t(np)$; (b) parallel efficiency $E(np) = \Delta t(1)/(\Delta t(np) \cdot np)$. The solid line indicates ideal speed-up/efficiency. ●, using an IBM SP2 system with nodes at 66 MHz clock speed and 128 MB core memory; ○, CRAY T3E LC 384 with DEC Alpha EV5.6 nodes at 600MHz clock speed and 512 MB core memory. Note that in '●' a 32 bit floating point representation is used, while '○' uses a 64 bit representation. Therefore the individual work load is heavier in the latter case.

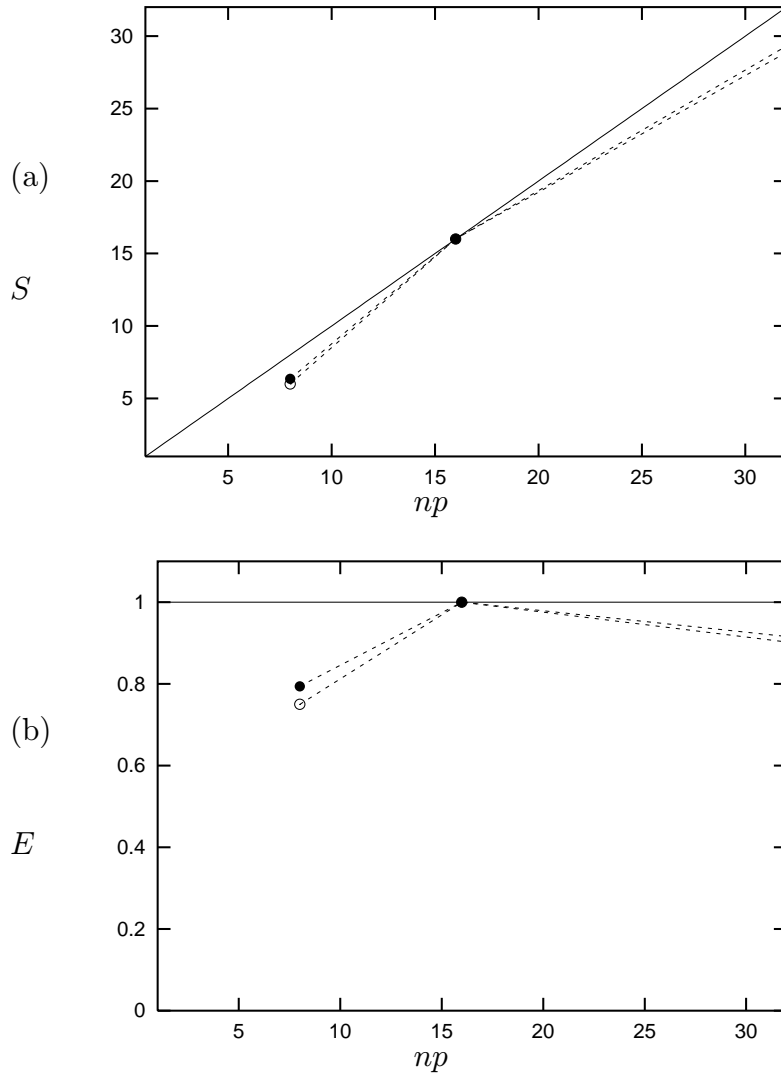


Figure 8: As figure 7 except that the problem size is 256^3 . Because of memory constraints, only computations with $n_P > 1$ are shown. The single execution time is approximated as : $t(1) = t(16) * 16$.

PIK Report-Reference:

- No. 1 3. Deutsche Klimatagung, Potsdam 11.-14. April 1994, Tagungsband der Vorträge und Poster (April 1994)
- No. 2 Extremer Nordsommer '92
Meteorologische Ausprägung, Wirkungen auf naturnahe und vom Menschen beeinflusste Ökosysteme, gesellschaftliche Perzeption und situationsbezogene politisch-administrative bzw. individuelle Maßnahmen (Vol. 1 - Vol. 4)
H.-J. Schellnhuber, W. Enke, M. Flechsig (Mai 1994)
- No. 3 Using Plant Functional Types in a Global Vegetation Model
W. Cramer (September 1994)
- No. 4 Interannual variability of Central European climate parameters and their relation to the large-scale circulation
P. C. Werner (Oktober 1994)
- No. 5 Coupling Global Models of Vegetation Structure and Ecosystem Processes - An Example from Arctic and Boreal Ecosystems
M. Plöchl, W. Cramer (Oktober 1994)
- No. 6 The use of a European forest model in North America: A study of ecosystem response to climate gradients
H. Bugmann, A. Solomon (Mai 1995)
- No. 7 A comparison of forest gap models: Model structure and behaviour
H. Bugmann, Y. Xiaodong, M. T. Sykes, Ph. Martin, M. Lindner, P. V. Desanker, S. G. Cumming (Mai 1995)
- No. 8 Simulating forest dynamics in complex topography using gridded climatic data
H. Bugmann, A. Fischlin (Mai 1995)
- No. 9 Application of two forest succession models at sites in Northeast Germany
P. Lasch, M. Lindner (Juni 1995)
- No. 10 Application of a forest succession model to a continentality gradient through Central Europe
M. Lindner, P. Lasch, W. Cramer (Juni 1995)
- No. 11 Possible Impacts of global warming on tundra and boreal forest ecosystems - Comparison of some biogeochemical models
M. Plöchl, W. Cramer (Juni 1995)
- No. 12 Wirkung von Klimaveränderungen auf Waldökosysteme
P. Lasch, M. Lindner (August 1995)
- No. 13 MOSES - Modellierung und Simulation ökologischer Systeme - Eine Sprachbeschreibung mit Anwendungsbeispielen
V. Wenzel, M. Kücken, M. Flechsig (Dezember 1995)
- No. 14 TOYS - Materials to the Brandenburg biosphere model / GAIA
Part 1 - Simple models of the "Climate + Biosphere" system
Yu. Svirezhev (ed.), A. Block, W. v. Bloh, V. Brovkin, A. Ganopolski, V. Petoukhov, V. Razzhevaikin (Januar 1996)
- No. 15 Änderung von Hochwassercharakteristiken im Zusammenhang mit Klimaänderungen - Stand der Forschung
A. Bronstert (April 1996)
- No. 16 Entwicklung eines Instruments zur Unterstützung der klimapolitischen Entscheidungsfindung
M. Leimbach (Mai 1996)
- No. 17 Hochwasser in Deutschland unter Aspekten globaler Veränderungen - Bericht über das DFG-Rundgespräch am 9. Oktober 1995 in Potsdam
A. Bronstert (ed.) (Juni 1996)
- No. 18 Integrated modelling of hydrology and water quality in mesoscale watersheds
V. Krysanova, D.-I. Müller-Wohlfeil, A. Becker (Juli 1996)
- No. 19 Identification of vulnerable subregions in the Elbe drainage basin under global change impact
V. Krysanova, D.-I. Müller-Wohlfeil, W. Cramer, A. Becker (Juli 1996)
- No. 20 Simulation of soil moisture patterns using a topography-based model at different scales
D.-I. Müller-Wohlfeil, W. Lahmer, W. Cramer, V. Krysanova (Juli 1996)
- No. 21 International relations and global climate change
D. Sprinz, U. Luterbacher (1st ed. July, 2nd ed. December 1996)
- No. 22 Modelling the possible impact of climate change on broad-scale vegetation structure - examples from Northern Europe
W. Cramer (August 1996)

- No. 23 A methode to estimate the statistical security for cluster separation
F.-W. Gerstengarbe, P.C. Werner (Oktober 1996)
- No. 24 Improving the behaviour of forest gap models along drought gradients
H. Bugmann, W. Cramer (Januar 1997)
- No. 25 The development of climate scenarios
P.C. Werner, F.-W. Gerstengarbe (Januar 1997)
- No. 26 On the Influence of Southern Hemisphere Winds on North Atlantic Deep Water Flow
S. Rahmstorf, M. H. England (Januar 1977)
- No. 27 Integrated systems analysis at PIK: A brief epistemology
A. Bronstert, V. Brovkin, M. Krol, M. Lüdeke, G. Petschel-Held, Yu. Svirezhev, V. Wenzel (März 1997)
- No. 28 Implementing carbon mitigation measures in the forestry sector - A review
M. Lindner (Mai 1997)
- No. 29 Implementation of a Parallel Version of a Regional Climate Model
M. Kücken, U. Schättler (Oktober 1997)
- No. 30 Comparing global models of terrestrial net primary productivity (NPP): Overview and key results
W. Cramer, D. W. Kicklighter, A. Bondeau, B. Moore III, G. Churkina, A. Ruimy, A. Schloss, participants of "Potsdam '95" (Oktober 1997)
- No. 31 Comparing global models of terrestrial net primary productivity (NPP): Analysis of the seasonal behaviour of NPP, LAI, FPAR along climatic gradients across ecotones
A. Bondeau, J. Kaduk, D. W. Kicklighter, participants of "Potsdam '95" (Oktober 1997)
- No. 32 Evaluation of the physiologically-based forest growth model FORSANA
R. Grote, M. Erhard, F. Suckow (November 1997)
- No. 33 Modelling the Global Carbon Cycle for the Past and Future Evolution of the Earth System
S. Franck, K. Kossacki, Ch. Bounama (Dezember 1997)
- No. 34 Simulation of the global bio-geophysical interactions during the Last Glacial Maximum
C. Kubatzki, M. Claussen (Januar 1998)
- No. 35 CLIMBER-2: A climate system model of intermediate complexity. Part I: Model description and performance for present climate
V. Petoukhov, A. Ganopolski, V. Brovkin, M. Claussen, A. Eliseev, C. Kubatzki, S. Rahmstorf (Februar 1998)
- No. 36 Geocybernetics: Controlling a rather complex dynamical system under uncertainty
H.-J. Schellnhuber, J. Kropp (Februar 1998)
- No. 37 Untersuchung der Auswirkungen erhöhter atmosphärischer CO₂-Konzentrationen auf Weizenbestände des Free-Air Carbondioxid Enrichment (FACE) - Experimentes Maricopa (USA)
Th. Kartschall, S. Grossman, P. Michaelis, F. Wechsung, J. Gräfe, K. Waloszczyk, G. Wechsung, E. Blum, M. Blum (Februar 1998)
- No. 38 Die Berücksichtigung natürlicher Störungen in der Vegetationsdynamik verschiedener Klimagebiete
K. Thonicke (Februar 1998)
- No. 39 Decadal Variability of the Thermohaline Ocean Circulation
S. Rahmstorf (März 1998)
- No. 40 SANA-Project results and PIK contributions
K. Bellmann, M. Erhard, M. Flechsig, R. Grote, F. Suckow (März 1998)
- No. 41 Umwelt und Sicherheit: Die Rolle von Umweltschwellenwerten in der empirisch-quantitativen Modellierung
D. F. Sprinz (März 1998)
- No. 42 Reversing Course: Germany's Response to the Challenge of Transboundary Air Pollution
D. F. Sprinz, A. Wahl (März 1998)
- No. 43 Modellierung des Wasser- und Stofftransportes in großen Einzugsgebieten. Zusammenstellung der Beiträge des Workshops am 15. Dezember 1997 in Potsdam
A. Bronstert, V. Krysanova, A. Schröder, A. Becker, H.-R. Bork (eds.) (April 1998)
- No. 44 Capabilities and Limitations of Physically Based Hydrological Modelling on the Hillslope Scale
A. Bronstert (April 1998)
- No. 45 Sensitivity Analysis of a Forest Gap Model Concerning Current and Future Climate Variability
P. Lasch, F. Suckow, G. Bürger, M. Lindner (Juli 1998)
- No. 46 Wirkung von Klimaveränderungen in mitteleuropäischen Wirtschaftswäldern
M. Lindner (Juli 1998)
- No. 47 SPRINT-S: A Parallelization Tool for Experiments with Simulation Models
M. Flechsig (Juli 1998)

- No. 48 The Odra/Oder Flood in Summer 1997: Proceedings of the European Expert Meeting in Potsdam, 18 May 1998
A. Bronstert, A. Ghazi, J. Hladny, Z. Kundzewicz, L. Menzel (eds.) (September 1998)
- No. 49 Struktur, Aufbau und statistische Programmbibliothek der meteorologischen Datenbank am Potsdam-Institut für Klimafolgenforschung
H. Österle, J. Glauer, M. Denhard (Januar 1999)
- No. 50 The complete non-hierarchical cluster analysis
F.-W. Gerstengarbe, P. C. Werner (Januar 1999)
- No. 51 Struktur der Amplitudengleichung des Klimas
A. Hauschild (April 1999)
- No. 52 Measuring the Effectiveness of International Environmental Regimes
C. Helm, D. F. Sprinz (Mai 1999)
- No. 53 Untersuchung der Auswirkungen erhöhter atmosphärischer CO₂-Konzentrationen innerhalb des Free-Air Carbon Dioxide Enrichment-Experimentes: Ableitung allgemeiner Modellösungen
Th. Kartschall, J. Gräfe, P. Michaelis, K. Waloszczyk, S. Grossman-Clarke (Juni 1999)
- No. 54 Flächenhafte Modellierung der Evapotranspiration mit TRAIN
L. Menzel (August 1999)
- No. 55 Dry atmosphere asymptotics
N. Botta, R. Klein, A. Almgren (September 1999)
- No. 56 Wachstum von Kiefern-Ökosystemen in Abhängigkeit von Klima und Stoffeintrag - Eine regionale Fallstudie auf Landschaftsebene
M. Erhard (Dezember 1999)
- No. 57 Response of a River Catchment to Climatic Change: Application of Expanded Downscaling to Northern Germany
D.-I. Müller-Wohlfeil, G. Bürger, W. Lahmer (Januar 2000)
- No. 58 Der "Index of Sustainable Economic Welfare" und die Neuen Bundesländer in der Übergangsphase
V. Wenzel, N. Herrmann (Februar 2000)
- No. 59 Weather Impacts on Natural, Social and Economic Systems (WISE, ENV4-CT97-0448) German report
M. Flechsig, K. Gerlinger, N. Herrmann, R. J. T. Klein, M. Schneider, H. Sterr, H.-J. Schellnhuber (Mai 2000)
- No. 60 The Need for De-Aliasing in a Chebyshev Pseudo-Spectral Method
M. Uhlmann (Juni 2000)
- No. 61 National and Regional Climate Change Impact Assessments in the Forestry Sector - Workshop Summary and Abstracts of Oral and Poster Presentations
M. Lindner (ed.) (Juli 2000)
- No. 62 Bewertung ausgewählter Waldfunktionen unter Klimaänderung in Brandenburg
A. Wenzel (August 2000)
- No. 63 Eine Methode zur Validierung von Klimamodellen für die Klimawirkungsforschung hinsichtlich der Wiedergabe extremer Ereignisse
U. Böhm (September 2000)
- No. 64 Die Wirkung von erhöhten atmosphärischen CO₂-Konzentrationen auf die Transpiration eines Weizenbestandes unter Berücksichtigung von Wasser- und Stickstofflimitierung
S. Grossman-Clarke (September 2000)
- No. 65 European Conference on Advances in Flood Research, Proceedings, (Vol. 1 - Vol. 2)
A. Bronstert, Ch. Bismuth, L. Menzel (eds.) (November 2000)
- No. 66 The Rising Tide of Green Unilateralism in World Trade Law - Options for Reconciling the Emerging North-South Conflict
F. Biermann (Dezember 2000)
- No. 67 Coupling Distributed Fortran Applications Using C++ Wrappers and the CORBA Sequence Type
Th. Slawig (Dezember 2000)
- No. 68 A Parallel Algorithm for the Discrete Orthogonal Wavelet Transform
M. Uhlmann (Dezember 2000)