

Universiteit Utrecht



Department
of
Mathematics

A simple and efficient parallel FFT algorithm
using the BSP model

by

Márcia A. Inda and Rob H. Bisseling

Preprint
nr. 1138
March 2000

A SIMPLE AND EFFICIENT PARALLEL FFT ALGORITHM USING THE BSP MODEL

MÁRCIA A. INDA AND ROB H. BISSELING

ABSTRACT. In this paper, we present a new parallel radix-4 FFT algorithm based on the BSP model. Our parallel algorithm uses the group-cyclic distribution family, which makes it simple to understand and easy to implement. We show how to reduce the communication cost of the algorithm by a factor of three, in the case that the input/output vector is in the cyclic distribution. We also show how to reduce computation time on computers with a cache-based architecture. We present performance results on a Cray T3E with up to 64 processors, obtaining reasonable efficiency levels for local problem sizes as small as 256 and very good efficiency levels for sizes larger than 2048.

1. INTRODUCTION

The *discrete Fourier transform* (DFT) plays an important role in computational science. DFT applications ranges from solving numerical differential equations to signal processing. (For an introduction to DFT applications see e.g. [7].) The widespread use of DFTs in computational science is mainly due to the existence of fast algorithms, known by the general name of fast Fourier transform (FFT), which compute the DFT of an input vector of size N in $O(N \log N)$ operations instead of the $O(N^2)$ operations needed by a direct approach, i.e., by a matrix-vector multiplication.

In 1965, Cooley and Tukey [9] published a paper describing the FFT idea (giving special attention to the so called *radix-2 FFT*). Since then, many variants of the algorithm have appeared. For an extensive discussion of the family of FFT algorithms, see Van Loan [27]. In recent years, after the dawn of parallel computing, the originally sequential FFT algorithms have been modified and adapted to the needs of parallel computation (see e.g. [1, 2, 3, 8, 11, 12, 14, 15, 21, 22, 25, 27]).

The lack of a unified parallel computing model and the existence of many different parallel architectures have made it rather difficult to develop efficient and portable parallel FFTs. Recently, however, as the parallel programming environments have become less machine dependent, examples of such algorithms have appeared. Typical examples are the

Key words and phrases. fast Fourier transform, bulk synchronous parallel, parallel computation.

Inda supported by a doctoral fellowship from CAPES, Brazil. Computer time on the Cray T3E of HPaC, Delft, funded by NCF, The Netherlands.

6-pass (or 6-step) approach and the related transpose approach (see e.g. [2, 3, 14, 15, 12]). Those algorithms regard the input vector of size $N = N_0N_1$ as an $N_0 \times N_1$ matrix, and carry out the computations in a similar way as done for two-dimensional FFTs. Since those algorithms require the number of processors p to be a divisor of N_0 and N_1 , they can only be used if $p \leq \sqrt{N}$.

As the number of available processors grows and the communication speed increases, it is important to develop parallel algorithms that can handle more than \sqrt{N} processors. Though generalized algorithms have already been proposed, they only work for very specific combinations of N and p such as $p = N^{(d-1)/d}/k$, where d is the dimensionality and $1 \leq k \leq \sqrt{N}$ [21, Chap. 10.3] or $k = 1$ [22], and both N and p are powers of two. Furthermore, to our knowledge none of those algorithms were implemented.

Our main aim in this paper is to present a new parallel FFT algorithm and its implementation. Our parallel algorithm works for any $p < N$ as long as both are powers of two, which is required because of the radix-2 framework. (A mixed-radix framework will be discussed elsewhere [18].) We dedicate the remainder of this section to giving a brief introduction to the basic framework of radix-2 and radix-4 FFTs and to the bulk synchronous parallel (BSP) model. In Section 2, we derive our parallel FFT algorithm by inserting suitable permutation matrices into the basic radix-2 decomposition of the Fourier matrix. This approach leads to a simple and easy to implement distributed memory parallel FFT algorithm. In Section 3, we present a set of templates that are used in the implementation of the algorithm. In Section 4, we present variants of our FFT algorithm. We show how to modify the algorithm to accept vectors that are not in the block distribution. We also show how to obtain a *cache-friendly* version of our algorithm, that is, an algorithm that takes advantage of the cache memory of a computer by breaking up the computations into small sections in such a way that the data stored in the cache is completely used before new data is brought in. In Section 5, we present results regarding the performance of our implementation and discuss aspects such as the *cache effect*. In Section 6, we draw our conclusions and discuss future work.

1.1. The fast Fourier transform. The DFT of a complex vector \mathbf{z} of size N is defined as the complex vector \mathbf{Z} , also of size N , with components

$$Z_k = \sum_{j=0}^{N-1} z_j e^{\frac{2\pi i j k}{N}}, \quad 0 \leq k < N. \quad (1.1)$$

The inverse DFT, which transforms the complex vector \mathbf{Z} back into the vector \mathbf{z} , is then defined by

$$z_j = \frac{1}{N} \sum_{k=0}^{N-1} Z_k e^{-\frac{2\pi i j k}{N}}, \quad 0 \leq j < N. \quad (1.2)$$

Alternatively, the DFT can be seen as a matrix-vector multiplication:

$$\mathbf{Z} = F_N \cdot \mathbf{z}. \quad (1.3)$$

The complex matrix F_N is known as the $N \times N$ *Fourier matrix*; it has elements $(F_N)_{jk} = w_N^{jk}$, where

$$w_N = e^{\frac{2\pi i}{N}}. \quad (1.4)$$

Though it is possible to develop FFT algorithms that compute the DFT of a vector of arbitrary size, the radix-2 FFT algorithm only works for N 's that are powers of two. For simplicity, we will restrict our discussion to such values of N . The sequential iterative radix-2 FFT algorithm starts with the so-called *bit reversal* permutation of the input vector (see Section 3.2), and proceeds in $\log_2 N$ butterfly stages, numbered $K = 2, 4, 8, \dots, N$, which modify the vector. Each butterfly stage consists of N/K times a butterfly computation:

$$\begin{pmatrix} z_{t+j} \\ z_{t+j+K/2} \end{pmatrix} \leftarrow \begin{pmatrix} z_{t+j} + w_K^j \cdot z_{t+j+K/2} \\ z_{t+j} - w_K^j \cdot z_{t+j+K/2} \end{pmatrix}, \quad \text{for } j = 0, 1, \dots, K/2 - 1, \quad (1.5)$$

where $t = 0, K, \dots, N - K$ indicates the beginning of a butterfly block in the vector. These operations cost one complex multiplication and two additions, or 10 real floating point operations (flops), per pair. The total flop count of the radix-2 FFT is therefore

$$C_{\text{FFT-2}}(N) = 10 \cdot \frac{K}{2} \cdot \frac{N}{K} \cdot \log_2 N = 5N \log_2 N.$$

Algorithm 1.1 is an in-place version of this algorithm.

Algorithm 1.1 Sequential radix-2 FFT algorithm.

CALL Seq_FFT(N, \mathbf{y}).**ARGUMENTS** N : Vector size; N is a power of 2 with $N \geq 2$. $\mathbf{y} = (y_0^{in}, \dots, y_{N-1}^{in})$: Complex vector of size N .**OUTPUT** $\mathbf{y} \leftarrow (y_0^{out}, \dots, y_{N-1}^{out})$, where $y_k^{out} = \sum_{j=0}^{N-1} y_j^{in} \exp(2\pi ijk/N)$.**DESCRIPTION**

1. Perform a bit reversal on \mathbf{y} .
2. Perform $\log_2 N$ butterfly stages $A_{K,N}$ on \mathbf{y} .

 $K \leftarrow 2$ **while** $K \leq N$ **do****for** $t = 0$ **to** $N - K$ **step** K **do****for** $j = 0$ **to** $K/2 - 1$ **do** $a \leftarrow w_K^j \cdot y_{t+j+K/2}$ $y_{t+j+K/2} \leftarrow y_{t+j} - a$ $y_{t+j} \leftarrow y_{t+j} + a$ $K \leftarrow 2 \cdot K$

Following van Loan's matrix approach [27], Algorithm 1.1 can be described as a sequence of sparse matrix-vector multiplications which correspond to the following decomposition of the Fourier matrix¹

$$F_N = A_{N,N} \cdots A_{8,N} A_{4,N} A_{2,N} P_N, \quad (1.6)$$

where P_N is the $N \times N$ permutation matrix corresponding to the bit reversal permutation (step 1 of Algorithm 1.1), and the $N \times N$ matrices $A_{K,N}$ correspond to the butterfly stages (step 2 of Algorithm 1.1). The block structure of the butterfly stages leads to block-diagonal matrices of the form

$$A_{K,N} = I_{N/K} \otimes B_K, \quad (1.7)$$

which is shorthand for a block-diagonal matrix $\text{diag}(B_K, \dots, B_K)$ with N/K copies of the $K \times K$ matrix B_K on the diagonal. The symbol \otimes represents the direct (or Kronecker) product of two matrices, which is formally defined at the end of this subsection. The matrix B_K is known as the $K \times K$ *2-butterfly matrix* which corresponds to the butterfly

¹Actually, the matrix decomposition corresponding to the algorithm of Cooley and Tukey [9] is $F_N = P_N \tilde{A}_{N,N} \cdots \tilde{A}_{8,N} \tilde{A}_{4,N} \tilde{A}_{2,N}$, where $\tilde{A}_{K,N} = P_N^{-1} A_{K,N} P_N$.

computation (1.5). This matrix can be written as

$$B_K = \begin{bmatrix} I_{K/2} & \Omega_{K/2} \\ I_{K/2} & -\Omega_{K/2} \end{bmatrix}. \quad (1.8)$$

Here, the matrix $I_{K/2}$ is the $K/2 \times K/2$ identity matrix and $\Omega_{K/2}$ is the $K/2 \times K/2$ diagonal matrix

$$\Omega_{K/2} = \text{diag}(w_K^0, w_K^1, \dots, w_K^{K/2-1}). \quad (1.9)$$

Later on we will also need generalized versions of $A_{K,N}$:

$$A_{K,N}^\alpha = I_{N/K} \otimes B_K^\alpha, \quad (1.10)$$

where B_K^α is the *generalized $K \times K$ 2-butterfly matrix* [3, 5, 10]

$$B_K^\alpha = \begin{bmatrix} I_{K/2} & \Omega_{K/2}^\alpha \\ I_{K/2} & -\Omega_{K/2}^\alpha \end{bmatrix}, \quad (1.11)$$

which has the same form as the original B_K , but the weights w_K^j in (1.9) are replaced by $w_K^{j+\alpha}$, where α can be any real number.

In practice, often a radix-4 FFT is used. A radix-4 algorithm can be derived completely analogously to the radix-2 algorithm, yielding a similar matrix decomposition. The algorithm starts with a reversal of pairs of bits instead of a reversal of single bits, and proceeds in $\log_4 N$ 4-butterfly stages which involve quadruples of vector components instead of pairs. Since 34 flops are performed per quadruple, this brings the flop count down to

$$C_{\text{FFT-4}}(N) = 34 \cdot \frac{K}{4} \cdot \frac{N}{K} \cdot \log_4 N = 4.25N \log_2 N.$$

The resulting algorithm has the disadvantage that either it must be assumed that N is a power of four, or special precautions must be taken which complicate the algorithm.

We take a slightly different approach: wherever possible we take pairs of stages $A_{K,N} A_{K/2,N}$ together and perform them as one operation. Our $K \times K$ 4-butterfly matrix has the form

$$D_K = B_K(I_2 \otimes B_{K/2}) = \begin{bmatrix} I_{K/4} & \Lambda_{K/4}^2 & \Lambda_{K/4} & \Lambda_{K/4}^3 \\ I_{K/4} & -\Lambda_{K/4}^2 & i\Lambda_{K/4} & -i\Lambda_{K/4}^3 \\ I_{K/4} & \Lambda_{K/4}^2 & -\Lambda_{K/4} & -\Lambda_{K/4}^3 \\ I_{K/4} & -\Lambda_{K/4}^2 & -i\Lambda_{K/4} & i\Lambda_{K/4}^3 \end{bmatrix}, \quad (1.12)$$

where $\Lambda_{K/4}$ is the $K/4 \times K/4$ diagonal matrix

$$\Lambda_{K/4} = \text{diag}(w_K^0, w_K^1, \dots, w_K^{K/4-1}). \quad (1.13)$$

This matrix is a symmetrically permuted version of the radix-4 butterfly matrix [27].² This approach gives the efficiency of a radix-4 FFT algorithm, and the flexibility of treating a parallel FFT within the radix-2 framework. For example, if we wish to permute the data sometime during the computation, for reasons of data locality, this can happen after any stage, and not only after an even number of stages.

An algorithm for the inverse FFT is obtained using the following property:

$$F_N^{-1} = \frac{1}{N} \bar{F}_N = \frac{1}{N} \bar{A}_{N,N} \cdots \bar{A}_{4,N} \bar{A}_{2,N} P_N. \quad (1.14)$$

The backward algorithm is basically the same as the forward one, the only difference being that the powers of w_K are replaced by their conjugates and that the final result is rescaled.

Now, we define the direct product of two matrices and give some properties that will be used in the course of the paper.

Definition 1.1 (Direct product). Let A be a $q \times r$ matrix and B be an $m \times n$ matrix. Then the direct product (or Kronecker product) of A and B is the $qm \times rn$ matrix defined by

$$A \otimes B = \begin{bmatrix} a_{0,0}B & \cdots & a_{0,r-1}B \\ \vdots & \ddots & \vdots \\ a_{q-1,0}B & \cdots & a_{q-1,r-1}B \end{bmatrix}.$$

As one would expect, the direct product is associative, but it is not commutative. Lemma 1.2 summarizes some direct product properties that follow directly from the definition. (See [24, 27] for other useful properties).

Lemma 1.2 (Properties of the direct product). *The following holds.*

1. $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$, as long as the products are defined.
2. $(I_m \otimes I_n) = I_{mn}$.

²In verifying this, note that van Loan defines the weights to be $w_K = \exp(-\frac{2\pi i}{K})$.

3. If A and B are square matrices of order m and n , respectively, then $(A \otimes I_n)(I_m \otimes B) = (A \otimes B) = (I_m \otimes B)(A \otimes I_n)$.
4. If A and B are square matrices of order n such that $AB = BA$, then $(I_m \otimes A)(I_m \otimes B) = (I_m \otimes B)(I_m \otimes A)$.

1.2. The bulk synchronous parallel model. The BSP model [26] is a parallel programming model which gives a simple and effective way to produce portable parallel algorithms. It does not depend on a specific computer architecture, and it provides a simple cost function that enables us to choose between algorithms without actually having to implement them. In the BSP model, a computer consists of a set of p processors, each with its own memory, connected by a communication network that allows processors to access the private memories of other processors. Accessing local memory (the processor's own memory) is faster than accessing remote memory (memory owned by other processors), but access time is considered to be independent from the computer architecture. In this model, algorithms consist of a sequence of *supersteps* and *synchronization barriers*. The use of supersteps and synchronization barriers imposes a sequential structure on parallel algorithms, and this greatly simplifies the design process.

The variant of the BSP model that we use is a *single program multiple data* (SPMD) model, i.e., each one of the p processors executes a copy of the same program, though each has its own data. The program distinguishes between the processors through a parameter s (the processor identification number). Special cases are treated using “if” statements. In our model, a superstep is either a *computation superstep*, or a *communication superstep*. A computation superstep is a sequence of local computations carried out on data already available locally before the start of the superstep. A communication superstep consists of communication of data between processors. To ensure the correct execution of the algorithm, global synchronization barriers (i.e., places of the algorithm where all processors must synchronize with each other) precede and/or follow a communication superstep.

A BSP computer can be characterized by four global parameters:

- p , the number of processors;
- v , the computing velocity in flop/s;

- g , the communication time per data element sent or received, measured in flop time units;
- l , the synchronization time, also measured in flop time units.

Algorithms can be analyzed by using the parameters p, g , and l . The parameter v is used to estimate the total execution time after the cost function has been computed. The flop count of a computation superstep is simply the maximum amount of work (in flops) of any processor. The flop count of a communication superstep is $hg + l$ or $hg + 2l$ (depending on the number of global synchronizations). Here h is the maximum number of data elements sent or received by any processor. The cost function of an algorithm can be obtained by adding the flops of the separate supersteps. This yields an expression of the form $a + bg + cl$. For further details and some basic techniques, see [4, 16]. The second reference describes BSplib, a standard library defined in May 1997 which enables parallel programming in BSP style. The Paderborn University BSP (PUB) library [6] is another library that permits programming in BSP style; it provides the extra feature of subset synchronization.

1.3. Parallel radix-2 FFTs. Since the introduction of parallel computers, and even before that, methods for parallelizing FFT algorithms have been proposed [23]. The earliest methods produced parallel algorithms that, using p processors, carry out an FFT of size N in $O(\log p)$ computation supersteps, which are interleaved by $O(\log p)$ communication supersteps that need to communicate $O(\frac{N}{p})$ data elements per processor (see e.g. [8, 11, 25]) and hence have cost $O(\frac{N}{p}g + l)$. Such methods appeared as a direct consequence of the divide-and-conquer structure of the radix-2 FFT algorithm. The paper [8] by Chu and George discusses several existing parallel algorithms of this type and three variants of their own. Restricting the discussion to vector sizes that are powers of two, they present a common framework in which all the algorithms they discuss are reorderings of one another in the following sense.

Each butterfly stage K of an FFT of size N , performs pairwise operations that combine elements j and $j + K/2$ from the vector being transformed using the weight $w_K^{j \bmod K}$. Writing j in its binary representation $j = (j_{m-1}, \dots, j_0)_2$, where $m = \log_2 N$, we observe that elements j and $j + K/2$ differ only in bit $\log_2 K - 1$ and that $w_K^{j \bmod K} = w_K^{(j_{\log_2 K-1}, \dots, j_0)_2}$. If the ordering of the vector is changed, so that original element j is stored as element

l , the butterfly stages must be modified to carry out the same operations. If we can represent the new ordering using a permutation of the original bits, it is easy to know which elements to combine and which weights to use. For example, if $N = 16$ a possible reordering of the input vector could be $l = (j_0, j_2, j_1, j_3)_2$, where $j = (j_3, j_2, j_1, j_0)_2$. The butterfly stage corresponding to $K = 16$ should then combine elements $l = (j_0, j_2, j_1, 0)_2$ with $l + 1 = (j_0, j_2, j_1, 1)_2$ using weights $w_{16}^{(j_2, j_1, j_0)_2}$.

In the parallel scenario, any group of $\log_2 p$ bits can be used to represent the processor number, while the remaining $\log_2(N/p)$ bits are used to represent the local index. If the bit corresponding to the next butterfly stage is one of the $\log_2(N/p)$ bits that represent the local index, then that stage is local, otherwise communication is needed.

Swarztrauber [25] carries out a similar discussion. He starts with a more general formulation of the problem, where N is not restricted to powers of two, but when discussing the distributed memory framework, he only considers FFTs on a hypercube, restricting both p and N to powers of two. The problem of the algorithms discussed in [8, 11, 25] is that reorderings are carried out by means of exchanging one bit at a time. Since there are $\log_2 p$ bits in the processor part, $\log_2 p$ communication supersteps of size $O(\frac{N}{p}g + l)$ are needed. A less expensive solution to the problem is to exchange all the processor bits with a group of local bits corresponding to butterfly stages that have already been performed. Since the communication cost of the permutation that exchanges many bits is of the same order $O(\frac{N}{p}g + l)$ as for exchanging one bit, the reduction in the communication cost is huge.

The basic idea for such algorithms already appears in the original Cooley and Tukey paper [9]. In their derivation of the FFT algorithm, they start by considering the case where N can be decomposed as $N = N_0 N_1$, and rewrite (1.1) as

$$Z_{k_1, k_0} = \sum_{j_1=0}^{N_1-1} \left(\sum_{j_0=0}^{N_0-1} z_{j_0, j_1} w_N^{j_0 k_0 N_1} \right) w_N^{j_1 (k_1 N_0 + k_0)}, \quad 0 \leq k_1 < N_1, 0 \leq k_0 < N_0, \quad (1.15)$$

where $Z_{k_1, k_0} = Z_{k_1 N_0 + k_0} = Z_k$, and $z_{j_0, j_1} = z_{j_0 N_1 + j_1} = z_j$. Since $w_N^{j_0 k_0 N_1} = w_{N_0}^{j_0 k_0}$, the inner sum of (1.15) corresponds to a DFT of size N_0 , which can be computed by the same process as before if N_0 is not prime. They remark that this process can be applied to any possible factorization of N , $N = N_0 \dots N_{H-1}$ and that, if N is composite enough, real gains (over the $O(N^2)$ direct approach) can be achieved. Afterwards they derive the

radix-2 algorithm by choosing N to be a power of two. If instead of decomposing N into its prime factors, we stop at a higher level, we obtain a decomposition of the FFT into a sequence of shorter FFTs that, in the parallel case, can be spread out over the processors. This is what happens in our FFT algorithm presented in the next section and in algorithms based on the 6-pass approach and the related transpose approach (see e.g. [2, 3, 14, 15, 12]).

2. THE PARALLEL ALGORITHM

2.1. Group-cyclic distribution family and the parallel FFT. Since our parallel FFT algorithm is based on the radix-2 decomposition (1.6) of the Fourier matrix, N must be a power of two. For practical reasons N/p must be integer and therefore p must also be a power of two. Our parallel FFT algorithm makes use of the data distributions defined below.

Definition 2.1 (Cyclic distribution in r groups, $C^r(p, N)$). Let r , p , and N be integers with $1 \leq r \leq p \leq N$, such that r divides p and N . Let \mathbf{f} be a vector of size N to be distributed over p processors organized in r groups. Define $M = N/r$ to be the size of the subvector of a group and $u = p/r$ to be the number of processors in a group. We say that \mathbf{f} is *cyclically distributed in r groups* (or *r -cyclically distributed*) *over p processors* if, for all j , the element f_j has local index $j' = (j \bmod M) \text{ div } u$, and is stored in processor $s_0 + s_1$, where $s_0 = (j \bmod M) \cdot u$ is the number of the first processor in the group (i.e., the processor offset) and $s_1 = (j \bmod M) \bmod u$ is the processor identification within the group.

We use the name *group-cyclic distribution family* to designate all the r -cyclic distributions generated by the same N and p . This family includes the well-known cyclic distribution ($C^1(p, N)$) and block distribution ($C^p(p, N)$) as extreme cases.

The parallel FFT algorithm works as follows. A total of $H = \lceil \log_2 N / \log_2(N/p) \rceil = \lceil \log_{\frac{N}{p}} N \rceil$ phases is performed. (The number of phases is the largest integer H for which $(N/p)^{H-1} \leq N$.) In phase 0, the algorithm uses the block distribution to perform the first $\log_2(N/p)$ butterfly stages, i.e., those involving butterflies with $K \leq N/p$ (which we call short distance butterflies). Afterwards, in each intermediate phase $1 \leq J < H - 1$ it uses the r -cyclic distribution, with $r = p/(N/p)^J$, to perform a group of $\log_2(N/p)$ butterfly




FIGURE 1. Butterfly operations using the group-cyclic distribution family $C^r(p, N) = C^r(8, 32)$. (A) Logical view, (B) storage view. The short distance butterflies ($A_{2,32}$ and $A_{4,32}$) are performed using the $C^8(8, 32)$ distribution (block distribution). The medium distance butterflies ($A_{8,32}$ and $A_{16,32}$) are performed using the $C^2(8, 32)$ distribution. The long distance butterflies ($A_{32,32}$) are performed using the $C^1(8, 32)$ distribution (cyclic distribution). For clarity, not all butterfly pairs are shown.

stages with $(N/p)^J < K \leq (N/p)^{J+1}$ (the medium distance butterflies). Finally, in phase $H - 1$, it uses the cyclic distribution to perform the remaining butterfly stages, i.e., those involving butterflies with $K > (N/p)^{H-1}$ (the long distance butterflies).

Figure 1 illustrates the use of the group-cyclic distribution family in our parallel FFT. The same operations are illustrated in two ways: (A) using the *logical view*, and (B) using the *storage view*. The logical view emphasizes the logical sequence of the elements in the vector while the storage view emphasizes the way the elements are actually stored. For the block distribution, both views are the same.

Our algorithm has only $O(\log N / \log(N/p))$ communication supersteps of size $O(\frac{N}{p}g + l)$, which is a significant improvement over the $O(\log p)$ communication supersteps also of

size $O(\frac{N}{p}g + l)$ of the algorithms discussed in the previous section. McColl [22] outlined a parallel FFT algorithm which is a special case of the FFT algorithm we present here. His algorithm only works for $N = (N/p)^H$. Furthermore, his algorithm sends the indices corresponding to the weights together with the data vector, increasing the communication costs unnecessarily. In the original BSP paper [26], Valiant already discussed the FFT and mentioned the order of the communication cost that can be achieved.

2.2. Permutations and permutation matrices. Let u , v , and N be positive integers such that u divides v and v divides N . We define the following permutation:

$$\begin{aligned}\gamma_{u,v,N} : \{0, \dots, N-1\} &\rightarrow \{0, \dots, N-1\} \\ j = j_0 \cdot M + j_1 \cdot u + j_2 &\mapsto l = j_0 \cdot M + j_2 \cdot \frac{N}{v} + j_1,\end{aligned}\tag{2.1}$$

where $M = \frac{N}{v}u$, $j_0 = j \text{ div } M$, $j_1 = (j \bmod M) \text{ div } u$, and $j_2 = j \bmod u$. Note that $\gamma_{u,v,N}^{-1} = \gamma_{\frac{N}{v}, \frac{N}{u}, N}$ and that $\gamma_{1,v,N} = \gamma_{u,N,N}$ is the identity permutation. Permuting a vector of size N by $\gamma_{u,v,N}$ can be achieved by dividing the vector into v/u subvectors of size $M = Nu/v$ and then performing a (perfect) shuffle permutation $\sigma_{u,M}$:

$$\begin{aligned}\sigma_{u,M} : \{0, \dots, M-1\} &\rightarrow \{0, \dots, M-1\} \\ j &\mapsto k = (j \bmod u) \cdot \frac{M}{u} + j \text{ div } u,\end{aligned}\tag{2.2}$$

on each of the subvectors.³ This relation can be expressed by

$$\gamma_{u,v,N}(j) = j \text{ div } M \cdot M + \sigma_{u,M}(j \bmod M),\tag{2.3}$$

which implies that $\gamma_{u,u,N} = \sigma_{u,N}$.

In the case that p divides N , redistributing a vector of size N from block distribution to r -cyclic distribution over p processors is equivalent to permuting it by $\gamma_{u,p,N}$, where $u = p/r$. Using matrix notation, this permutation can be expressed by the $N \times N$ permutation matrix:

$$(\Gamma_{u,p,N})_{lj} = \begin{cases} 1, & \text{if } l = \gamma_{u,p,N}(j), \\ 0, & \text{otherwise.} \end{cases}\tag{2.4}$$

³The shuffle permutations $\sigma_{p,N}(j)$ and $\sigma_{\frac{N}{p},N}(j)$ can be used to permute a vector of size N distributed over p processors from block to cyclic distribution and vice-versa.

Multiplying a vector \mathbf{y} by $\Gamma_{u,p,N}$ results in a vector with components $(\Gamma_{u,p,N}\mathbf{y})_l = y_{\gamma_{u,p,N}^{-1}(l)}$, for all l ; in other words, this multiplication corresponds to redistributing the vector from block distribution to cyclic distribution in $r = p/u$ groups. Note that $\Gamma_{u,p,N} = I_r \otimes S_{\frac{N}{p}u}$, cf. (2.3). The matrix corresponding to the inverse permutation $\gamma_{u,p,N}^{-1}$ is $\Gamma_{u,p,N}^{-1} = \Gamma_{u,p,N}^T = \Gamma_{\frac{N}{p},\frac{N}{u},N}$. From now on, we use the abbreviations Γ_u and γ_u to denote $\Gamma_{u,p,N}$ and $\gamma_{u,p,N}$, respectively. We restrict the use of subscripts p and N to cases where they are not obvious from the context. We also use $S_{u,M}$ to denote $\Gamma_{u,u,M}$.

2.3. Decomposition of the Fourier matrix. To obtain the parallel FFT algorithm, we modify the original radix-2 decomposition of the Fourier matrix (1.6) by inserting identity permutation matrices $I_N = \Gamma_u^{-1}\Gamma_u$ corresponding to the changes of distribution, and regrouping the matrices in the resulting decomposition. In the case that $p \leq \sqrt{N}$ this is done as follows:

$$\begin{aligned} F_N &= \Gamma_p^{-1}\Gamma_p \cdot A_{N,N} \cdot \Gamma_p^{-1}\Gamma_p \cdots \Gamma_p^{-1}\Gamma_p \cdot A_{2\frac{N}{p},N} \cdot \Gamma_p^{-1}\Gamma_p \cdot A_{\frac{N}{p},N} \cdots A_{2,N}P_N \\ &= \Gamma_p^{-1}(\Gamma_p A_{N,N} \Gamma_p^{-1}) \cdots (\Gamma_p A_{2\frac{N}{p},N} \Gamma_p^{-1}) \Gamma_p \cdot A_{\frac{N}{p},N} \cdots A_{2,N}P_N. \end{aligned} \quad (2.5)$$

By defining

$$\hat{A}_{k,u,p,N} = \Gamma_{u,p,N} A_{ku,N} \Gamma_{u,p,N}^{-1}, \quad (2.6)$$

and using the fact that Γ_1 is the identity matrix, we can rewrite (2.5) as

$$F_N = \Gamma_p^{-1} \cdot \underbrace{\hat{A}_{\frac{N}{p},p,p,N} \cdots \hat{A}_{2\frac{N}{p^2},p,p,N}}_{\text{phase 1}} \cdot \Gamma_p \Gamma_1^{-1} \cdot \underbrace{\hat{A}_{\frac{N}{p},1,p,N} \cdots \hat{A}_{2,1,p,N}}_{\text{phase 0}} \cdot \Gamma_1 P_N. \quad (2.7)$$

As we did with the permutation matrices Γ , from now on we denote $\hat{A}_{k,u,p,N}$ by $\hat{A}_{k,u}$, reserving the indices p and N for when they are not obvious from the context and for stand-alone definitions. In the general case, not restricted to $p \leq \sqrt{N}$, we arrive at the following decomposition of the Fourier matrix:

$$\begin{aligned} F_N &= \Gamma_p^{-1} \underbrace{\hat{A}_{\frac{N}{p},p} \cdots \hat{A}_{2\frac{(N/p)^{H-1}}{p},p}}_{\text{phase } H-1} \Gamma_p \cdot \Gamma_{(\frac{N}{p})^{H-2}}^{-1} \underbrace{\hat{A}_{\frac{N}{p},(\frac{N}{p})^{H-2}} \cdots \hat{A}_{2,(\frac{N}{p})^{H-2}}}_{\text{phase } H-2} \Gamma_{(\frac{N}{p})^{H-2}} \cdots \\ &\quad \cdots \Gamma_{\frac{N}{p}}^{-1} \underbrace{\hat{A}_{\frac{N}{p},\frac{N}{p}} \cdots \hat{A}_{2,\frac{N}{p}}}_{\text{phase 1}} \Gamma_{\frac{N}{p}} \cdot \Gamma_1^{-1} \underbrace{\hat{A}_{\frac{N}{p},1} \cdots \hat{A}_{2,1}}_{\text{phase 0}} \Gamma_1 \cdot P_N. \end{aligned} \quad (2.8)$$




FIGURE 2. (A) Structure of the $N \times N$ matrix $\hat{A}_{k,u} = I_r \otimes \text{diag}(A_{k,n}^{0/u}, A_{k,n}^{1/u}, \dots, A_{k,n}^{(u-1)/u})$. Example with $N = 128$, $p = 8$, $r = 2$, $k = 8$, and hence $u = 4$ and $n = 16$. (For clarity, the $A_{k,n}^\alpha$ are depicted as A_k^α) (B) Matrix $\text{diag}(A_{k,n}^{0/u}, A_{k,n}^{1/u}, \dots, A_{k,n}^{(u-1)/u})$. (C) Matrices $B_k^{s_1/u}$, with $s_1 = 0, \dots, u - 1$.

The matrices $\hat{A}_{k,u}$ are block diagonal matrices with block size $\frac{N}{p}$:

$$\hat{A}_{k,u,p,N} = I_r \otimes \text{diag}(A_{k,n}^{0/u}, A_{k,n}^{1/u}, \dots, A_{k,n}^{(u-1)/u}), \quad (2.9)$$

where $r = p/u$, $n = N/p$, and $A_{k,n}^\alpha$ was defined previously, cf. (1.10). Figure 2 exemplifies the structure of the matrix $\hat{A}_{k,u}$. We shall formally state this as Corollary 2.3 which follows from Theorem 2.2.

Theorem 2.2. *Let u , M , and k be powers of two such that $2 \leq k \leq M/u$. Define $n = M/u$. Then*

$$S_{u,M} A_{ku,M} S_{u,M}^{-1} = \text{diag}(A_{k,n}^{0/u}, A_{k,n}^{1/u}, \dots, A_{k,n}^{(u-1)/u}).$$

Proof. See Appendix A. □

Corollary 2.3. *Let r , p , and N be powers of two with $1 \leq r \leq p < N$. Define $u = p/r$ and $n = N/p$. Let k be a power of two with $2 \leq k \leq n$. Then*

$$\hat{A}_{k,u,p,N} = I_r \otimes \text{diag}(A_{k,n}^{0/u}, A_{k,n}^{1/u}, \dots, A_{k,n}^{(u-1)/u}).$$

Proof. Define $M = N/r$. Then

$$\begin{aligned}\hat{A}_{k,u,p,N} &= \Gamma_{u,p,N} A_{ku,N} \Gamma_{u,p,N}^{-1} = (I_r \otimes S_{u,M})(I_r \otimes A_{ku,M})(I_r \otimes S_{u,M}^{-1}) \\ &= I_r \otimes (S_{u,M} A_{ku,M} S_{u,M}^{-1}) = I_r \otimes \text{diag}(A_{k,n}^{0/u}, A_{k,n}^{1/u}, \dots, A_{k,n}^{(u-1)/u}).\end{aligned}$$

□

Starting from the Fourier matrix decomposition (2.8), it is easy to develop a parallel (BSP) FFT algorithm. Since all the matrices $\hat{A}_{k,u}$ are block diagonal matrices with block size equal to N/p , any multiplication $\hat{A}_{k,u} \cdot \mathbf{y}$ can be handled locally, provided that the vector \mathbf{y} is in the block distribution. This property guarantees that matrix decomposition (2.8) detaches communication and computation completely. On the one hand, each generalized butterfly phase $(\hat{A}_{\frac{N}{p},u} \dots \hat{A}_{2,u}) \cdot \mathbf{y}$, is a strict computation superstep. On the other hand, each permutation Γ_u is a strict communication superstep. In the next section we give a complete description of the resulting parallel algorithm.

3. IMPLEMENTATION OF THE PARALLEL ALGORITHM

We present our algorithm in the form of *templates*, which are high level of detail algorithms, ready to be implemented, though not necessarily completely optimized. In the list that follows we introduce the terminology used in our parallel algorithms.

- **Supersteps.** Each superstep is numbered textually and labeled according to its type: (**Comp**) computation superstep, (**Comm**) communication superstep, (**CpCm**) subroutine containing both computation and communication supersteps. Global synchronizations are explicitly indicated by the keyword *Synchronize*. Supersteps inside loops are executed repeatedly, though they are numbered only once.
- **Indexing.** All the indices of vectors or array structures are global. This means that array elements have a unique index which is independent of the processor that owns it. This property enables us to describe variables and gain access to arrays in an unambiguous manner, even though the array is distributed and each processor has

only part of it. (In an actual implementation, it is more convenient to convert the indexing scheme to a local one.)

- **Communication.** Communication between processors is indicated using

$$\mathbf{g}_j \leftarrow \text{Put}(pid, n, \mathbf{f}_i)$$

This operation puts n elements of array \mathbf{f} , starting from element i , into processor pid and stores them there in array \mathbf{g} starting from element j . Subscripts are not needed when the first element of the array is 0 or when communicating scalars. When communicating more than one element, we use boldface to emphasize that we are dealing with a vector and not a scalar. All puts are assumed to be buffered, so that they are safely carried out, even if \mathbf{f} and \mathbf{g} happen to be the same.

Algorithm 3.1 is a direct implementation of the matrix decomposition (2.8). The input vector \mathbf{y} is transformed in place. Superstep 1 permutes the input vector by a bit reversal: $\mathbf{y} \leftarrow P_N \cdot \mathbf{y}$. Superstep 2 carries out the short distance butterflies: $\mathbf{y} \leftarrow (A_{\frac{N}{p}, N} \dots A_{2, N}) \cdot \mathbf{y}$. (Since γ_1 is the identity permutation, $\hat{A}_{k,1} = A_{k,N}$, for $k = 2, \dots, N/p$.) Superstep 3 permutes \mathbf{y} to the r -cyclic distribution, with $r = \max(1, p/(N/p))$: $\mathbf{y} \leftarrow \Gamma_{p/r} \cdot \mathbf{y}$. Each time superstep 4 is executed, it computes a group of medium distance butterflies: $\mathbf{y} \leftarrow (\hat{A}_{\frac{N}{p}, (\frac{N}{p})^J} \dots \hat{A}_{2, (\frac{N}{p})^J}) \cdot \mathbf{y}$, $1 \leq J \leq H - 2$. Superstep 5 prepares the vector for the next butterfly phase by permuting it to the r -cyclic distribution, with $r = \max(1, p/(N/p)^{J+1})$: $\mathbf{y} \leftarrow (\Gamma_r \Gamma_{(\frac{N}{p})^J}^{-1}) \cdot \mathbf{y}$. Superstep 6 carries out the long distance butterflies: $\mathbf{y} \leftarrow (\hat{A}_{\frac{N}{p}, p} \dots \hat{A}_{2, \frac{(N/p)^{H-1}}{p}, p}) \cdot \mathbf{y}$. Finally, superstep 7 permutes the vector back to the block distribution: $\mathbf{y} \leftarrow \Gamma_p^{-1} \cdot \mathbf{y}$. Note that, to obtain the normalized inverse transform, the output vector must be divided by N . The subroutines used in the FFT template are described in the following subsections.

3.1. Generalized butterflies. The subroutine BTFLY (Algorithm 3.2) is a sequential subroutine that multiplies the input vector by $A_{n,n}^\alpha \dots A_{k_0,n}^\alpha A_{k_0/2,n}^\alpha$. Step 1 performs pairs of generalized butterfly operations. The k -th iteration of the outermost while-loop performs the pair of generalized butterflies stages $A_{k,n}^\alpha A_{k/2,n}^\alpha$, the t -th iteration of the intermediate for-loop corresponds to the t -th repetition of the *generalized 4-butterfly* $D_k^\alpha = B_k^\alpha(I_2 \otimes B_{k/2}^\alpha)$, cf. (1.12), which is computed by the innermost for-loop. Step 2

Algorithm 3.1 Template for the parallel fast Fourier transform, using the group-cyclic distribution family.

CALL `BSP_FFT($s, p, sign, N, \mathbf{y}$)`.

ARGUMENTS

- s : Processor identification; $0 \leq s < p$.
- p : Number of processors; p is a power of 2 with $p < N$.
- $sign$: Transform direction; +1 for forward, -1 for backward.
- N : Transform size; N is a power of 2 with $N \geq 2$.
- $\mathbf{y} = (y_0^{in}, \dots, y_{N-1}^{in})$: Complex vector of size N (block distributed).

OUTPUT $\mathbf{y} \leftarrow (y_0^{out}, \dots, y_{N-1}^{out})$, where $y_k^{out} = \sum_{j=0}^{N-1} y_j^{in} \exp(sign \cdot 2\pi i j k / N)$.

DESCRIPTION

1^{CpCm}	Parallel bit reversal permutation. <code>BSP_BitRev(s, p, N, \mathbf{y})</code>
2^{Comp}	Short distance butterflies. <code>BTFLY($0, sign, \frac{N}{p}, 4, \mathbf{y}_{s \frac{N}{p}}$)</code>
3^{CpCm}	Permutation to $C^{\max(1, p/(N/p))}(p, N)$ distribution. <code>BSP_BlockToCyclic($s - s \bmod \frac{N}{p}, s \bmod \frac{N}{p}, \min(p, \frac{N}{p}), \frac{N}{p}, \mathbf{y}_{(s-s \bmod \frac{N}{p}) \frac{N}{p}}$)</code>
	$H \leftarrow \lceil \log_{\frac{N}{p}} N \rceil$
	for $J = 1$ to $H - 2$ do
4^{Comp}	Medium distance butterflies. <code>BTFLY($\frac{s \bmod (N/p)^J}{(N/p)^J}, sign, \frac{N}{p}, 4, \mathbf{y}_{s \frac{N}{p}}$)</code>
5^{Comm}	Permutation to $C^{\max(1, p/(N/p)^{J+1})}(p, N)$ distribution. <code>BSP_CyclicToCyclic($s, p, N, (N/p)^J, \min(p, (N/p)^{J+1}), \mathbf{y}$)</code>
6^{Comp}	Long distance butterflies. <code>BTFLY($\frac{s}{p}, sign, \frac{N}{p}, 4 \frac{(N/p)^{H-1}}{p}, \mathbf{y}_{s \frac{N}{p}}$)</code>
7^{CpCm}	Permutation to block distribution. <code>BSP_CyclicToBlock($0, s, p, \frac{N}{p}, \mathbf{y}$)</code>

is only executed if the number of butterfly stages is odd. It computes the last generalized 2-butterfly $A_{n,n}^\alpha$. The FFT algorithm computes the desired (short, medium, or long distance) butterfly stages corresponding to phase J , $0 \leq J < H$, by defining the input parameter $\alpha = (s \bmod u)/u$, where $u = \min(p, (N/p)^J)$, and performing the generalized butterfly stages on the local part of the vector \mathbf{y} (i.e., the subvector $\mathbf{y}_{s \frac{N}{p}}$ of size N/p that starts at element sN/p).

If the needed weights are stored in a lookup table, the cost of an $A_{k,n}^\alpha A_{k/2,n}^\alpha$ butterfly operation is $34 \cdot \frac{k}{4} \cdot \frac{n}{k} = \frac{17}{2}n$. Summing over all pairs $A_{k,n}^\alpha A_{k/2,n}^\alpha$ and adding 5n flops for

Algorithm 3.2 Template for the sequential generalized butterfly operations.

CALL BTFLY($\alpha, sign, n, k_0, \mathbf{y}$).**ARGUMENTS**

α : Butterfly parameter, used to compute the correct weights; $0 \leq \alpha < 1$.
 $sign$: Transform direction; +1 for forward, -1 for backward.

n : Vector size; n is a power of 2 with $n \geq 2$.

k_0 : Smaller 4-butterfly size; k_0 is a power of 2 with $4 \leq k_0 \leq 2n$.

$\mathbf{y} = (y_0, \dots, y_{n-1})$: Complex vector of size n .

OUTPUT

$\mathbf{y} \leftarrow A_{n,n}^\alpha \dots A_{k_0,n}^\alpha A_{k_0/2,n}^\alpha \mathbf{y}$ for forward, and
 $\mathbf{y} \leftarrow \bar{A}_{n,n}^\alpha \dots \bar{A}_{k_0,n}^\alpha \bar{A}_{k_0/2,n}^\alpha \mathbf{y}$ for backward.

DESCRIPTION

1. Perform pairs of butterfly stages $A_{k,n}^\alpha A_{k/2,n}^\alpha$.
 $k \leftarrow k_0$
while $k \leq n$ **do**
 for $t = 0$ **to** $n - k$ **step** k **do**
 for $j = 0$ **to** $k/4 - 1$ **do**
 $yw1 \leftarrow w_k^{sign \cdot (j+\alpha)} \cdot y_{t+j+k/2}$
 $yw2 \leftarrow w_k^{sign \cdot 2(j+\alpha)} \cdot y_{t+j+k/4}$
 $yw3 \leftarrow w_k^{sign \cdot 3(j+\alpha)} \cdot y_{t+j+3k/4}$
 $a \leftarrow y_{t+j} + yw2$
 $b \leftarrow y_{t+j} - yw2$
 $c \leftarrow yw1 + yw3$
 $d \leftarrow yw1 - yw3$
 $y_{t+j} \leftarrow a + c$
 $y_{t+j+k/4} \leftarrow b + sign \cdot id$
 $y_{t+j+k/2} \leftarrow a - c$
 $y_{t+j+3k/4} \leftarrow b - sign \cdot id$
 $k \leftarrow 4 \cdot k$
 2. Perform the last butterfly stage $A_{n,n}^\alpha$.
if $k = 2n$ **then**
 for $j = 0$ **to** $n/2 - 1$ **do**
 $a \leftarrow w_n^{sign \cdot (j+\alpha)} \cdot y_{j+n/2}$
 $y_{j+n/2} \leftarrow y_j - a$
 $y_j \leftarrow y_j + a$
-

the last butterfly (if necessary) gives a cost of

$$\begin{aligned} C_{\text{BTFLY}}(n, k_0) &= \\ &= \frac{17}{2}n \cdot [(\log_2 n - \log_2 k_0 + 2) \text{ div } 2] + 5n \cdot [(\log_2 n - \log_2 k_0 + 2) \bmod 2] \quad (3.1) \\ &= \frac{17}{4}n \cdot [\log_2 n - \log_2 k_0 + 2] + \frac{3}{4}n \cdot [(\log_2 n - \log_2 k_0 + 2) \bmod 2], \end{aligned}$$

for the generalized butterfly algorithm (Algorithm 3.2).

The total computation cost of our parallel FFT (Algorithm 3.1) is obtained by adding the costs of the butterfly phases:

$$C_{\text{BTFLY}}\left(\frac{N}{p}, 4\right) = \frac{17}{4} \frac{N}{p} \log_2 \frac{N}{p} + \frac{3}{4} \frac{N}{p} (\log_2 \frac{N}{p} \bmod 2),$$

for phases $J = 0$ to $H - 2$, where $H = \lceil \log_{\frac{N}{p}} N \rceil$, and for the last phase $H - 1$ if $(N/p)^{H-1} = p$, i.e., $H = \log_{\frac{N}{p}} N$. Otherwise, the cost for the last phase is

$$\begin{aligned} C_{\text{BTFLY}}\left(\frac{N}{p}, 4 \frac{(N/p)^{H-1}}{p}\right) &= \frac{17}{4} \frac{N}{p} (\log_2 N \bmod \log_2 \frac{N}{p}) \\ &\quad + \frac{3}{4} \frac{N}{p} [(\log_2 N \bmod \log_2 \frac{N}{p}) \bmod 2]. \end{aligned}$$

This gives

$$\begin{aligned} C_{\text{FFT,par,comp}}(N, p) &= \frac{17}{4} \frac{N}{p} \log_2 N + \frac{3}{4} \frac{N}{p} [(\log_2 \frac{N}{p} \bmod 2)(\log_2 N \bmod \log_2 \frac{N}{p}) \\ &\quad + (\log_2 N \bmod \log_2 \frac{N}{p}) \bmod 2], \end{aligned} \tag{3.2}$$

where the second term corresponds to the extra cost we have to pay for performing 2-butterflies. The communication and synchronization costs of our parallel FFT are discussed in Section 3.5 after we discuss the parallel permutation subroutines.

3.2. Parallel bit reversal. The bit reversal matrix P_N is defined by

$$(P_N)_{jk} = \begin{cases} 1, & \text{if } j = \text{rev}_N(k), \\ 0, & \text{otherwise.} \end{cases} \tag{3.3}$$

Here, rev_N is the bit reversal permutation

$$\begin{aligned} \text{rev}_N : \{0, \dots, N - 1\} &\rightarrow \{0, \dots, N - 1\} \\ j &= \sum_{l=0}^{m-1} b_l 2^l \mapsto k = \sum_{l=0}^{m-1} b_{m-l-1} 2^l, \end{aligned} \tag{3.4}$$

where $m = \log_2 N$ and $(b_{m-1} \dots b_0)_2$ is the binary representation of j . Note that $\text{rev}_N^{-1} = \text{rev}_N$, which means that $P_N^{-1} = P_N$.

The bit reversal permutation has the following very useful property.

Lemma 3.1. *Let $u = 2^q$, $N = 2^m$, with $q \leq m$. Then*

$$\text{rev}_N(j) = \text{rev}_{\frac{N}{u}}(j \bmod u) + \frac{N}{u} \cdot \text{rev}_u(j \bmod u), \quad 0 \leq j < N.$$

Proof. Let $d = m - q$, so that $\frac{N}{u} = 2^d$. If the binary representation of j is $(b_{m-1} \dots b_0)_2$, then

$$j = j \bmod u + u \cdot (j \text{ div } u) = \sum_{l=0}^{q-1} b_l 2^l + 2^q \sum_{l=0}^{d-1} b_{l+q} 2^l.$$

Now,

$$\text{rev}_N(j) = \sum_{l=0}^{m-1} b_{m-l-1} 2^l = \sum_{l=0}^{d-1} b_{m-l-1} 2^l + 2^d \cdot \sum_{l=0}^{q-1} b_{q-l-1} 2^l = a + \frac{N}{u} \cdot b.$$

If we show that $a = \text{rev}_{\frac{N}{u}}(j \text{ div } u)$ and $b = \text{rev}_u(j \bmod u)$, we are done. In fact,

$$\begin{aligned} a &= \sum_{l=0}^{d-1} b_{m-l-1} 2^l = \sum_{l=0}^{d-1} b_{m-q-l-1+q} 2^l = \quad (\text{substituting } c_l = b_{l+q}) \\ &= \sum_{l=0}^{d-1} c_{d-l-1} 2^l = \text{rev}_{\frac{N}{u}} \left(\sum_{l=0}^{d-1} c_l 2^l \right) = \text{rev}_{\frac{N}{u}} \left(\sum_{l=0}^{d-1} b_{l+q} 2^l \right) = \text{rev}_{\frac{N}{u}}(j \text{ div } u) \end{aligned}$$

and

$$b = \sum_{l=0}^{q-1} b_{q-l-1} 2^l = \text{rev}_u \left(\sum_{l=0}^{q-1} b_l 2^l \right) = \text{rev}_u(j \bmod u).$$

□

Corollary 3.2. *Let $u \leq N$ be powers of two. Then*

$$P_N = (I_u \otimes P_{\frac{N}{u}})(P_u \otimes I_{\frac{N}{u}})S_{u,N}$$

Proof. The matrix $(I_u \otimes P_{\frac{N}{u}})(P_u \otimes I_{\frac{N}{u}})S_{u,N}$ corresponds to a sequence of three permutations:

1. $j \rightarrow l = \sigma_{u,N}(j) = j \bmod u \cdot \frac{N}{u} + j \text{ div } u$
2. $l \rightarrow t = \text{rev}_u(l \text{ div } \frac{N}{u}) \cdot \frac{N}{u} + l \bmod \frac{N}{u} = \text{rev}_u(j \bmod u) \cdot \frac{N}{u} + j \text{ div } u$
3. $t \rightarrow k = t \text{ div } \frac{N}{u} \cdot \frac{N}{u} + \text{rev}_{\frac{N}{u}}(t \bmod \frac{N}{u}) = \text{rev}_u(j \bmod u) \cdot \frac{N}{u} + \text{rev}_{\frac{N}{u}}(j \text{ div } u) = \text{rev}_N(j)$

□

Let \mathbf{y} be a vector of size $N = 2^m$ block distributed over $p = 2^q$ processors. Suppose that we want to permute it by a bit reversal permutation, i.e., perform $\mathbf{y} \leftarrow P_N \cdot \mathbf{y}$. Applying Corollary 3.2 with $u = p$, it is possible to split the parallel bit reversal permutation in two parts as follows.

1. $\mathbf{y} \leftarrow (P_p \otimes I_{\frac{N}{p}}) S_{p,N} \cdot \mathbf{y}$, which is a global permutation that sends the elements to the final destination processors, but with local indices still in the original order:

$$j \rightarrow t = \underbrace{\text{rev}_p(j \bmod p)}_{\text{Proc}(t)} \cdot \frac{N}{p} + \underbrace{j \div p}_{t'}$$

Having as basis the block distribution, from now on, we use $\text{Proc}(k) = k \div \frac{N}{p}$ to denote the processor in which element k is stored, and $k' = k \bmod \frac{N}{p}$ to denote the local index of the element.

2. $\mathbf{y} \leftarrow (I_p \otimes P_{\frac{N}{p}}) \cdot \mathbf{y}$, which is a local bit reversal permutation in the local index t' :

$$t' \rightarrow k' = \text{rev}_{\frac{N}{p}}(t')$$

Algorithm 3.3 carries out a parallel bit reversal using this idea. If we combine the local bit reversal (superstep 2 of Algorithm 3.3) with the short distance butterfly phase (superstep 2 of Algorithm 3.1), we obtain a complete local sequential FFT. This means that we can easily replace the two supersteps by any optimized FFT subroutine we can lay our hands on.

Algorithm 3.3 Template for the parallel bit reversal.

CALL `BSP_BitRev(s, p, n, \mathbf{y})`.

ARGUMENTS

- s : Processor identification; $0 \leq s < p$.
- p : Number of processors; p is a power of 2 with $p < N$.
- N : Vector size; N is a power of 2 with $N \geq 2$.
- $\mathbf{y} = (y_0, \dots, y_{N-1})$: Complex vector of size N (block distributed).

OUTPUT $\mathbf{y} \leftarrow P_N \mathbf{y}$.

DESCRIPTION

```

1Comm   Global permutation.
        for  $j = s \frac{N}{p}$  to  $s \frac{N}{p} + \frac{N}{p} - 1$  do
            dest  $\leftarrow \text{rev}_p(j \bmod p)$ 
             $x_{\text{dest}, \frac{N}{p} + j \div p} \leftarrow \text{Put}(\text{dest}, 1, y_j)$ 
        Synchronize
2Comp   Local bit reversal.
        for  $t' = 0$  to  $\frac{N}{p} - 1$  do
             $y_{s \frac{N}{p} + \text{rev}_{\frac{N}{p}}(t')} \leftarrow x_{s \frac{N}{p} + t'}$ 

```

If $p < N/p$, it is possible to optimize communication superstep 1 of Algorithm 3.3 by sending packets of data. This is done in a similar way as when permuting from block to

cyclic distribution (see Section 3.4); the only difference is in the destination processor, which is $\text{rev}_p(j \bmod p)$ instead of $j \bmod p$.

3.3. Permutations within the group-cyclic distribution family. Permuting a vector from the $C^{r_1}(p, N)$ distribution to the $C^{r_2}(p, N)$ distribution, where $r_1 = p/u_1$ and $r_2 = p/u_2$ may be any possible group sizes, not necessarily powers of two, can be done as follows: first, use $\gamma_{u_1}^{-1}$ to permute the vector to the block distribution, and then use γ_{u_2} to permute it to the $C^{r_2}(p, N)$ distribution. This operation is expensive if performed in parallel, because all the data have to be moved twice around the processors. The best approach is to combine both permutations into one:

$$\begin{aligned}\gamma_{u_1}^{u_2} : \{0, \dots, N-1\} &\rightarrow \{0, \dots, N-1\} \\ j &\mapsto l = \gamma_{u_2}(\gamma_{u_1}^{-1}(j)).\end{aligned}\tag{3.5}$$

(Note that $(\gamma_{u_1}^{u_2})^{-1} = \gamma_{u_2}^{u_1}$, and that $\gamma_{u_1}^{u_2}$ is an abbreviation for $\gamma_{u_1, p, N}^{u_2}$.) In the general case, there is no simple formula for computing the destination index l . Algorithm 3.4 is a template for this case. Some combinations of the parameters r_1, r_2, p , and N , however, lead to simpler expressions for the destination index. The simplest case is when r_1 or r_2 is equal to p , i.e., one of the distributions involved is the block distribution. This situation occurs in supersteps 3 and 7 of the FFT algorithm (Algorithm 3.1) and is discussed in Section 3.4.

3.4. Permutation from block to cyclic distribution. The permutations $\sigma_{p, N}$ and $\sigma_{p, N}^{-1}$ are the permutations that convert a vector from block to cyclic distribution and vice versa. In the case that $p < b = N/p$, both $\sigma_{p, N}$ and $\sigma_{p, N}^{-1}$ can be optimized by sending packets of size b/p (here we assume that p divides b , but nothing more).

For $\sigma_{p, N}$, this is done as follows. First perform a local permutation $\sigma_{p, b}$ on the local index j' ,

$$j' \rightarrow t' = j' \bmod p \cdot \frac{b}{p} + j' \text{ div } p.$$

Then perform a global *cyclic permutation of packets* on the global index $t = j \text{ div } b \cdot b + t' = t_0 \cdot b + t_1 \cdot \frac{b}{p} + t_2$,

$$t \rightarrow k = \underbrace{t_1}_{\text{Proc}(k)} \cdot b + \underbrace{t_0 \cdot \frac{b}{p}}_{k'} + t_2.\tag{3.6}$$

Algorithm 3.4 Template for the parallel permutation from r_1 -cyclic to r_2 -cyclic distribution.

CALL `BSP_CyclicToCyclic($s, p, N, u_1, u_2, \mathbf{y}$)`.

ARGUMENTS

s : Processor identification; $0 \leq s < p$.

p : Number of processors; $p < N$.

N : Vector size.

u_1 : Number of processors in the old group; $u_1 = p/r_1$.

u_2 : Number of processors in the new group; $u_2 = p/r_2$.

$\mathbf{y} = (y_0, \dots, y_{N-1})$: Complex vector of size N (block distributed).

OUTPUT $\mathbf{y} \leftarrow \Gamma_{u_2} \Gamma_{u_1}^{-1} \cdot \mathbf{y}$.

DESCRIPTION

```
1comm   Global permutation  $\gamma_{u_1}^{u_2}$ .
        for  $j = s\frac{N}{p}$  to  $s\frac{N}{p} + \frac{N}{p} - 1$  do
             $l \leftarrow \gamma_{u_1}^{u_2}(j)$ 
             $proc \leftarrow l \text{ div } \frac{N}{p}$ 
             $y_l \leftarrow \text{Put}(proc, 1, y_j)$ 
        Synchronize
```

The resulting global index k is then

$$\begin{aligned} k &= j' \bmod p \cdot b + j \text{ div } b \cdot \frac{b}{p} + j' \text{ div } p \\ &= j \bmod p \cdot b + j \text{ div } b \cdot \frac{b}{p} + (j \bmod b) \text{ div } p \\ &= j \bmod p \cdot b + (j \text{ div } b \cdot b + j \bmod b) \text{ div } p \\ &= j \bmod p \cdot b + j \text{ div } p, \end{aligned}$$

which is indeed $\sigma_{p,N}(j)$, see Figure 3.

Subroutine `BSP_BlockToCyclic` (Algorithm 3.5) with $s_1 = s$ and $s_0 = 0$ carries out the permutation $\sigma_{p,N}$. The subroutine performs the permutation on a vector \mathbf{y} of size $N = pb$ which is block distributed over p processors. This is done in superstep 3 of the FFT algorithm if $p \leq N/p$.

Subroutine `BSP_BlockToCyclic` can also be used to carry out the permutation $\gamma_{p, rp, rN}$, which is needed in superstep 3 of the FFT algorithm if $p > N/p$. This is possible because permuting a vector of size rN by $\gamma_{p, rp, rN}$ is the same as dividing it into r subvectors of size N and then performing a shuffle permutation $\sigma_{p,N}$ on each of the subvectors, cf. (2.3). In this case $s_1 = s \bmod p$, $s_0 = s - s_1$, and \mathbf{y} is a subvector starting at element s_0b of the original vector.




FIGURE 3. Schematic representation of a two stage permutation from block to cyclic distribution (storage view). Example with $N = 32$ and $p = 4$. (A) Local $\sigma_{4,8}$ permutation. (B) Global cyclic permutation of packets of size 2.

The template for subroutine `BSP_CyclicToBlock`, which carries out $\sigma_{p,N}^{-1}$ and $\gamma_{p,rp,rN}^{-1}$ is obtained in a similar way as the template for subroutine `BSP_BlockToCyclic`. The template starts by performing a global cyclic permutation of packets, and then it carries out a local permutation $\sigma_{p,b}^{-1}$. (The template is not presented here; see [17] for more details.)

3.5. BSP cost. To compute the total cost of our parallel FFT algorithm (Algorithm 3.1) we need to sum the computation, communication, and synchronization costs. The computation costs were already obtained in Section 3.1. To simplify the final result we only include the higher order term of the total computation cost (3.2),

$$C_{\text{FFT,par,comp}}(N, p) = \frac{17}{4} \frac{N}{p} \log_2 N, \quad (3.7)$$

which is exact when only 4-butterflies are performed.

The communication and synchronization costs are the costs involved in performing the bit reversal and the permutations related to the group-cyclic distribution family. The maximum amount of data sent or received during a permutation involving complex numbers is equal to N/p complex values (or $2N/p$ real values). If the permutation is performed with puts, the number of synchronizations is 1, giving a total cost of

$$C_{\text{permute}}(N, p) = 2 \frac{N}{p} \cdot g + l \quad (3.8)$$

Algorithm 3.5 Template for the parallel permutation from block to cyclic distribution.

CALL `BSP_BlockToCyclic($s_0, s_1, p, b, \mathbf{y}$).`
ARGUMENTS
 s_0, s_1 : Processor offset and processor identification within group; $0 \leq s_1 < p$.

 p : Number of processors in group.

 b : Block size; p divides b , if $p < b$.

 $\mathbf{y} = (y_0, \dots, y_{pb-1})$: Complex vector of size pb (block distributed within group).

OUTPUT $\mathbf{y} \leftarrow S_{p,pb}\mathbf{y}$.

DESCRIPTION

```

1comm   if  $p \geq b$  then
           Global  $\sigma_{p,pb}$  permutation.
           for  $j = s_1 \cdot b$  to  $(s_1 + 1) \cdot b - 1$  do
                $dest \leftarrow j \bmod p$ 
                $y_{dest \cdot b + j \bmod p} \leftarrow \text{Put}(s_0 + dest, 1, y_j)$ 
           Synchronize
2comp   else
           Local  $\sigma_{p,b}$  permutation.
           for  $j' = 0$  to  $b - 1$  do
                $k' \leftarrow j' \bmod p \cdot \frac{b}{p} + j' \bmod p$ 
                $x_{s_1 \cdot b + k'} \leftarrow y_{s_1 \cdot b + j'}$ 
3comm   Global cyclic permutation of packets.
           for  $proc = 0$  to  $p - 1$  do
                $y_{proc \cdot b + s_1 \cdot \frac{b}{p}} \leftarrow \text{Put}(s_0 + proc, \frac{b}{p}, x_{s_1 \cdot b + proc \cdot \frac{b}{p}})$ 
           Synchronize

```

for each of the $\lceil \log_{\frac{N}{p}} N \rceil + 1$ permutations performed in the FFT algorithm. The total cost of the FFT algorithm is

$$C_{\text{FFT, par}}(N, p) = \frac{17}{4} \frac{N}{p} \log_2 N + 2 \frac{N}{p} (\lceil \log_{\frac{N}{p}} N \rceil + 1) \cdot g + (\lceil \log_{\frac{N}{p}} N \rceil + 1) \cdot l. \quad (3.9)$$

In Section 5, we discuss the validity of cost function (3.9) as an accurate estimator of the true cost of the FFT algorithm.

4. VARIANTS OF THE ALGORITHM

4.1. Parallel FFT using other data distributions. Up to now, we discussed an FFT algorithm where the input and output (I/O) vector must be block distributed. There exist applications of the FFT where it would be better if the I/O vector would be distributed by other distributions or where the distribution can be freely chosen (see e.g. [14, 19, 28]). Here, we discuss how to modify our parallel FFT algorithm to accept I/O vectors that are not distributed by the block distribution.

The first and the last supersteps of Algorithm 3.1 are permutations. Because of this, the algorithm can be modified to accept any I/O data distribution without any extra communication cost, or even at a smaller communication cost depending on the desired distributions. If the input vector is not in the block distribution, the algorithm is modified by combining the redistribution to block distribution with the bit reversal permutation. If the output vector is expected to be in a distribution other than the block distribution, this is done by replacing the permutation from cyclic to block distribution by a permutation from the cyclic to the desired distribution.

If the desired distribution for the output vector is the cyclic distribution, the last communication superstep can be completely skipped. The first permutation can also be skipped if the input vector is already stored by the distribution associated with the bit reversal permutation. Applications where the input vector is bit reversed and the output vector is cyclically distributed are advantageous, because, in such cases, two complete permutations can be skipped.⁴ This saves two thirds of the total communication cost in the common case that $p \leq N/p$, leaving only one permutation in the middle of the computation.

While the cyclic distribution is simple and widely used, the distribution associated with the bit reversal permutation is awkward. Fortunately, it is possible to modify Algorithm 3.1 so that the cyclic distribution is a natural input distribution, i.e., a distribution that does not involve any communication as the first superstep. This is done as follows. The first three supersteps of Algorithm 3.1 are described by the matrix decomposition

$$\Gamma_u \cdot A_{\frac{N}{p}, N} \dots A_{2, N} \cdot P_N, \quad (4.1)$$

where $u = \min(p, N/p)$. Knowing that $A_{K,N} = I_p \otimes A_{K,\frac{N}{p}}$, and that the bit reversal matrix can be decomposed as $P_N = (I_p \otimes P_{\frac{N}{p}}) \cdot (P_p \otimes I_{\frac{N}{p}}) \cdot S_{p,N}$ (cf. Corollary 3.2), we rewrite

⁴The idea of skipping permutations to save communication time or to reduce the overhead caused by local permutations is known. Cooley and Tukey [9] already suggested this in order to save local bit reversals. Other authors (e.g., [14, 19, 28]) give examples where skipping permutations saves communication time.

matrix (4.1) as

$$\begin{aligned} & \Gamma_u \cdot (I_p \otimes A_{\frac{N}{p}, \frac{N}{p}}) \cdots (I_p \otimes A_{2, \frac{N}{p}}) \cdot (I_p \otimes P_{\frac{N}{p}}) \cdot (P_p \otimes I_{\frac{N}{p}}) \cdot S_{p,N} \\ &= \Gamma_u \cdot (I_p \otimes F_{\frac{N}{p}}) \cdot (P_p \otimes I_{\frac{N}{p}}) \cdot S_{p,N} = \Gamma_u \cdot (P_p \otimes I_{\frac{N}{p}}) \cdot (I_p \otimes F_{\frac{N}{p}}) \cdot S_{p,N}. \end{aligned} \quad (4.2)$$

Here we used Lemma 1.2. The first three supersteps of the parallel FFT algorithm derived from this new decomposition are: (1^{comm}) permutation from block to cyclic distribution, (2^{comp}) local FFT, (3^{comm}) permutation defined by $\Gamma_{\min(p, N/p)} \cdot (P_p \otimes I_{\frac{N}{p}})$. In the case that the input vector is already cyclically distributed, the first superstep can be skipped.

4.2. Generalized butterfly phase with adjustable size. In our original algorithm, we chose to insert the permutation matrices Γ_u in the leftmost possible position. This procedure corresponds to factoring N as $N = \frac{N}{(N/p)^{H-1}}(N/p)^{H-1}$, and gives an algorithm with a minimum number of permutations. However, if $p \neq (N/p)^{H-1}$, it is possible to insert the permutation matrices at an earlier position without increasing the number of permutations. The resulting algorithm would correspond to a different factorization of N .

We can use this flexibility to reduce the computation cost of some combinations of p and N by inserting the permutations so that a maximal number of generalized butterfly stages are paired off. Another reason to permute the vector at an earlier stage is that the sizes of the butterfly phases can be better balanced (so that all factors of N have approximately the same size). This would enhance the performance on a cache-sensitive computer (see the discussion in Section 5). An even a more effective way of enhancing the performance on a cache-sensitive computer is to reduce the butterfly sizes so that they always fit completely into the cache. We suggest a method in the following subsection.

4.3. Cache-friendly parallel FFT. Each computation superstep of our parallel FFT algorithm performs a butterfly phase which consists of a sequence of generalized butterfly stages represented by the operation $\mathbf{y} \leftarrow R_{l,n}^\alpha \mathbf{y}$, where l and n are powers of two with $2 \leq l \leq n$, and

$$R_{l,n}^\alpha = A_{n,n}^\alpha \cdots A_{2l,n}^\alpha A_n^\alpha, \quad (4.3)$$

is an $n \times n$ matrix. Suppose that the cache memory of a computer is such that the data needed by a butterfly phase of size n/v , where $v < n$ is a power of two, fits totally in

the computer cache. We can view v as the number of virtual processors available in each processor. If we decompose (4.3) into a sequence of smaller butterfly phases of size less than or equal to n/v which can be carried out independently from each other, we can get the most out of the cache of the computer.

Define $h = \lceil \log_{\frac{n}{v}} n \rceil$ and $j = \lceil \log_{\frac{n}{v}} l \rceil - 1$, so that $(\frac{n}{v})^j < l \leq (\frac{n}{v})^{j+1}$. Similarly to (2.8), if we denote $\Gamma_{u,v,n}$ by Γ_u , we can write

$$R_{l,n}^\alpha = \Gamma_v^{-1} \underbrace{\hat{A}_{\frac{n}{v},v}^\alpha \dots \hat{A}_{2(\frac{n}{v})^{h-1},v}^\alpha}_{\text{phase } h-j-1} \Gamma_v \cdot \Gamma_{(\frac{n}{v})^{h-2}}^{-1} \underbrace{\hat{A}_{\frac{n}{v},(\frac{n}{v})^{h-2}}^\alpha \dots \hat{A}_{2,(\frac{n}{v})^{h-2}}^\alpha}_{\text{phase } h-j-2} \Gamma_{(\frac{n}{v})^{h-2}} \dots \dots \dots \Gamma_{(\frac{n}{v})^{j+1}}^{-1} \underbrace{\hat{A}_{\frac{n}{v},(\frac{n}{v})^{j+1}}^\alpha \dots \hat{A}_{2,(\frac{n}{v})^{j+1}}^\alpha}_{\text{phase 1}} \Gamma_{(\frac{n}{v})^{j+1}} \cdot \Gamma_{(\frac{n}{v})^j}^{-1} \underbrace{\hat{A}_{\frac{n}{v},(\frac{n}{v})^j}^\alpha \dots \hat{A}_{\frac{l}{(\frac{n}{v})^j},(\frac{n}{v})^j}^\alpha}_{\text{phase 0}} \Gamma_{(\frac{n}{v})^j}, \quad (4.4)$$

where the $n \times n$ matrix $\hat{A}_{k,u}^\alpha$ is an abbreviation for $\hat{A}_{k,u,v,n}^\alpha = \Gamma_{u,v,n} A_{ku,n}^\alpha \Gamma_{u,v,n}^{-1}$. Generalized versions of Theorem 2.2 and Corollary 2.3 can be used to prove that

$$\hat{A}_{k,u,v,n}^\alpha = I_{\frac{n}{u}} \otimes \text{diag}(A_{k,\frac{n}{u}}^{\alpha/u}, A_{k,\frac{n}{u}}^{(\alpha+1)/u}, \dots, A_{k,\frac{n}{u}}^{(\alpha+u-1)/u}). \quad (4.5)$$

The matrix decomposition (4.4) can be used to construct an alternative (cache-friendly) algorithm for the computation of the generalized butterfly phases. Note that if $\alpha = 0$ then the resulting algorithm can be used to construct a cache-friendly sequential FFT algorithm.

5. EXPERIMENTAL RESULTS AND DISCUSSION

In this section, we present results on the performance of our implementation of the FFT. We implemented the FFT algorithm for the block distribution in ANSI C using the BSPlib communications library [16]. Our programs are completely self-contained, and we did not rely on any system-provided numerical software such as BLAS, FFTs, etc.

We tested our implementation on a Cray T3E with up to 64 processors, each having a theoretical peak speed of 600 Mflop/s. The accuracy of double precision (64-bit) arithmetic is 1.0×10^{-15} . We also give accuracy results from calculations on a Sun Workstation using IEEE 754 floating point arithmetic, which has a double precision accuracy of 2.2×10^{-16} , and which is the standard used in many computers such as workstations. To make a consistent comparison of the results, we compiled all test programs using the bspfront driver with options `-O3 -flibrary-level 2 -fcombine-puts` and measured

the elapsed execution times on exclusively dedicated CPUs using the system clock. The times given correspond to an average of the execution times of a forward FFT and a normalized backward FFT.

5.1. Accuracy. We tested the overall accuracy of our implementation by measuring the error obtained when transforming a random complex vector \mathbf{f} with values $\text{Re}(f_j)$ and $\text{Im}(f_j)$ uniformly distributed between 0 and 1. The relative error is defined as $\|\mathbf{F}^* - \mathbf{F}\|_2 / \|\mathbf{F}\|_2$, where \mathbf{F}^* is the vector obtained by transforming the original vector \mathbf{f} by a forward (or backward) FFT, and \mathbf{F} is the exact transform, which we computed using the same algorithm but using quadruple precision. Here, $\|\cdot\|_2$ indicates the L^2 -norm.

Table 1 shows the relative errors of the sequential algorithm for various problem sizes. Since the error for the forward and backward FFT are approximately the same, we present only the results for the forward transform. The errors of the parallel implementation are of the same order as in the sequential case. In fact, the error of the parallel implementation only differs from the error of the sequential one if the butterfly stages are not paired in the same way. This validates the parallel algorithm. The results also indicate that IEEE arithmetic is superior to the CRAY-specific arithmetic.

TABLE 1. Relative errors for the sequential FFT algorithm.

N	CRAY T3E	IEEE 754
512	2.4×10^{-16}	1.9×10^{-16}
1024	5.2×10^{-16}	1.6×10^{-16}
2048	8.4×10^{-16}	1.8×10^{-16}
4096	2.1×10^{-15}	1.9×10^{-16}
8192	3.2×10^{-15}	2.0×10^{-16}
16384	6.5×10^{-15}	2.2×10^{-16}
32768	2.3×10^{-14}	2.3×10^{-16}
65536	3.4×10^{-14}	2.3×10^{-16}

5.2. Performance of the sequential implementation. Our sequential FFT algorithm was implemented using Algorithm 3.2 with $\alpha = 0$. Its efficiency can be analyzed by looking at its execution times or its *FFT flop rates*:

$$\text{FFT}^{\text{rate}}(\text{seq}, N) = \frac{5N \log_2 N}{\text{Time}(\text{seq}, N)}, \quad (5.1)$$

where $\text{Time}(seq, N)$ is the execution time of the sequential implementation. Analyzing the performance of an FFT algorithm by using the number of flops of the radix-2 FFT as basis is a standard and useful procedure. By doing so, it is possible to compare different algorithms with different cost functions and also evaluate the overall performance of the algorithm as a function of N .

TABLE 2. Timing results (in ms) and FFT flop rates (in Mflop/s) of the sequential FFT on the Cray T3E.

N	Time	FFT ^{rate}
32	0.02	33.4
64	0.03	56.9
128	0.10	47.3
256	0.15	69.6
512	0.41	56.5
1024	0.66	77.5
2048	1.82	62.1
4096	2.94	83.5
8192	19.95	26.7
16384	58.91	19.5
32768	149.79	16.4
65536	318.28	16.5

Table 2 gives timing results and FFT flop rates for various problem sizes. The flop rates show that the performance of the algorithm increases until $N = 4096$, when it suddenly drops. This sudden decrease in performance happens because the data space allocated by the program becomes too large to fit completely in the cache memory of the CRAY T3E,⁵ which means that the computation becomes more expensive, because more accesses to the main memory are needed.

5.3. Scalability of the parallel implementation. The timing results obtained by our parallel algorithm are summarized in Table 3. We also present the theoretical predictions using the cost function (3.9) and the values of the BSP parameters v , g , and l listed in

⁵The cache size of the CRAY T3E is 96 Kbytes, which means that a sequential FFT of size up to $N = 4096$ fits completely in the cache (64 Kbytes for the data vector + 8 Kbytes for the weights table).

TABLE 3. Predicted and measured execution times (in ms) for the FFT on a Cray T3E. Boldface entries indicate out-of-cache computations

p	512		1024		2048		4096	
	pred	exp	pred	exp	pred	exp	pred	exp
seq	0.56	0.41	1.25	0.66	2.74	1.82	5.99	2.94
1	0.58	0.40	1.28	0.66	2.81	1.81	6.12	2.95
2	0.37	0.42	0.77	0.90	1.61	1.66	3.44	3.90
4	0.25	0.28	0.45	0.47	0.89	0.99	1.83	1.81
8	0.21	0.21	0.32	0.33	0.56	0.69	1.06	1.22
16	0.20	0.22	0.25	0.26	0.37	0.39	0.63	0.56
32	0.26	0.29	0.23	0.33	0.29	0.37	0.42	0.53
64	0.46	0.31	0.48	0.38	0.51	0.55	0.46	0.63
p	8192		16384		32768		65536	
	pred	exp	pred	exp	pred	exp	pred	exp
seq	12.97	19.95	27.93	58.91	59.9	149.8	127.7	318.3
1	13.23	19.95	28.46	58.98	60.9	149.8	129.8	315.8
2	7.33	8.93	15.61	33.65	33.2	87.1	70.3	207.3
4	3.83	4.40	8.09	9.72	17.1	39.7	36.1	101.1
8	2.12	2.33	4.36	5.28	9.1	12.5	19.1	46.7
16	1.16	1.20	2.30	2.26	4.7	5.4	9.8	12.7
32	0.70	0.75	1.29	1.43	2.5	2.9	5.1	7.1
64	0.61	0.76	0.91	0.98	1.5	1.7	2.9	3.2

Table 4 which were obtained using a modified version of the benchmark program from BSPpack⁶ [17].

Except for the out-of-cache computations (boldface entries in the table), the timings show that the BSP cost function predicts the behavior of the parallel implementation well. The discrepancy between predicted and measured results for out-of-cache computations is to be expected, since the computation speed, which we assumed to be constant, suddenly drops when the computations cannot be done completely in cache. These results show that the BSP model is a valid tool for analyzing and predicting parallel performance.

⁶Available at <http://www.math.uu.nl/people/bisseling/software.html>

TABLE 4. BSP parameters for the CRAY T3E, with $v = 34.9$ Mflop/s. The value of v is based on in-cache dot product computations.

p	g		l	
	(flops)	(μ s)	(flops)	(μ s)
1	0.28	0.00804	3	0.09
2	1.14	0.03279	479	13.72
4	1.46	0.04175	858	24.57
8	2.14	0.06133	1377	39.48
16	2.30	0.06600	1754	50.26
32	2.77	0.07940	2024	58.00
64	3.05	0.08758	3861	110.88

A way of analyzing the scalability of a parallel implementation is to look at its *absolute efficiency*

$$E^{abs}(p, N) = \frac{\text{Time}(\text{seq}, N)}{p\text{Time}(p, N)}, \quad (5.2)$$

as done in Figure 4. In theory, $E^{abs}(p, N) \leq 1$, and our goal is to achieve efficiencies as close to 1 as possible. The figure shows moderate efficiencies for small problem sizes ($N \leq 4096$). For $N \geq 8192$, efficiencies of up to 1.7 are achieved. Such amazing efficiencies are possible because of the so called *cache effect*: when $N \geq 8192$ the total amount of memory needed by the FFT is too large to fit in the cache memory of one processor, but, if the problem is executed using a sufficiently large number of processors, the memory required by each processor becomes small enough to fit in the cache. This effect is welcome, but it masks the real scalability of the algorithm.

Note that there is a sudden rise in the flop rate when the local problem size becomes small enough to fit in the cache. In this way the cache effect can be easily spotted and the scalability of the algorithm better judged. FFT sizes that fit completely in the cache ($N \leq 4096$) have a completely different behavior than larger problems. For small sizes ($N \leq 4096$) the efficiency decreases notably in going from one to two processors, then it is more or less constant up to 8–16 processors and after that it decreases steadily. For large sizes ($N > 16384$) the flop rate is nearly constant, both before and after the transition out-of-cache/in-cache, indicating a good scalability. The cases $8192 \leq N \leq 16384$ are intermediate cases where there is an increase in the efficiency in going from one to two




FIGURE 4. Absolute efficiencies of the FFT on a Cray T3E.

processors, but a deterioration of performance when the number of processors becomes too large.

We can also examine the scalability of our parallel algorithm by increasing the problem size together with the number of processors [13, 21], for instance by maintaining the local problem size N/p constant and increasing p . In doing so, we can learn about the asymptotic behavior of our algorithm. Figure 5 shows the predicted and measured efficiencies as a function of p for various values of N/p . The predicted values converge to a horizontal line as N/p increases which means that asymptotically the efficiency can be maintained at a constant level if N/p is maintained constant. The experimental results must be analyzed keeping in mind the cache effect, which causes the sudden increase in the efficiency. It is clear that efficiency can be maintained at reasonable levels for N/p as small as 256, and at very good levels for $N/p = 4096$.




FIGURE 5. Efficiency as a function of p for a constant problem size N/p on a Cray T3E. Dashed lines: predicted values. Solid lines: measured values.

6. CONCLUSIONS AND FUTURE WORK

In Section 2, we presented a new parallel FFT algorithm, Algorithm 3.1. This algorithm is a mixed radix-2 and radix-4 FFT. It was derived from the matrix decomposition corresponding to the radix-2 algorithm by inserting suitable permutation matrices corresponding to the group-cyclic distribution family. The use of the group-cyclic distribution family gives a parallel algorithm which is simple to understand and easy to implement.

The use of matrix notation proved to be a powerful tool for deriving and adapting the parallel FFT algorithms to our needs. With the help of matrix notation, we showed how to modify our original algorithm to accept I/O vectors that are not in block distribution, without incurring extra communication cost. Indeed, if the vector is cyclically distributed, we showed how to eliminate the first and the last permutation altogether, reducing the communication to one third of the original cost. Since the cyclic distribution is simple and widely used, this property can be exploited to obtain faster applications. A prime application of the cyclic distribution would be in the field of quantum molecular dynamics where a potential energy operator is applied to the input vector representing a wave packet and a kinetic energy operator is applied to the output vector [20]. Since both operations

are componentwise, any distribution can be chosen and hence the cyclic distribution is preferred.

We presented results concerning the performance of our implementation of Algorithm 3.1. The tests were carried out on a Cray T3E with up to 64 processors. Our implementation proved to scale reasonably well for small problem sizes ($N \leq 4096$) with up to 8 processors, and to scale very well for larger problem sizes ($N \geq 16384$). In part, the very favorable results obtained for larger N are due to the cache effect. Because of this effect, we analyzed our results in terms of FFT flop rate per processor, and also by using the theoretical cost function. Both analyses confirmed the previous results.

We also analyzed the asymptotic behavior of our algorithm by maintaining the local problem size N/p constant and increasing p . We concluded that the efficiency level is maintained as long as N/p is large enough. A study of the experimental data obtained on the Cray T3E indicates that reasonable efficiency levels ($E(p, N) \approx 0.5$) are already maintained for N/p as small as 256, and good efficiency levels are maintained for $N/p > 4096$.

Because the cache-based architecture of the Cray T3E influences our results so much, and there are many other computers with a similar architecture, we proposed the use of cache-friendly FFT algorithms. A cache-friendly sequential algorithm can be derived from our parallel algorithm by substituting the processors by virtual processors. It is also possible to derive a cache-friendly parallel algorithm by writing each generalized butterfly phase as a sequence of smaller generalized butterfly phases. We expect the scalability of such an algorithm to be similar to the theoretical scalability of our algorithm.

APPENDIX A. PROOF OF THEOREM 2.2

The proof of Theorem 2.2 uses the following lemma.

Lemma A.1. *Let u , M , and k be powers of two such that $u \leq M$ and $2 \leq k \leq M/u$. Define $K = ku$. Let j be an index, $0 \leq j < M$. Then*

1. *If $j \bmod K < K/2$, then $\sigma_{u,M}(j) \bmod k < k/2$.*
2. *If $j + K/2 < M$, then $\sigma_{u,M}(j + K/2) = \sigma_{u,M}(j) + k/2$.*

3. If $j_1 = j \bmod \frac{M}{u}$, and $j_0 = j \operatorname{div} \frac{M}{u}$, then

$$\frac{\sigma_{u,M}^{-1}(j) \bmod K}{K} = \frac{j_1 \bmod k + j_0/u}{k}.$$

Proof. Part 1: $\sigma_{u,M}(j) \bmod k = (j \bmod u \cdot \frac{M}{u} + j \operatorname{div} u) \bmod k = (j \operatorname{div} u) \bmod k$. Now, $j \operatorname{div} u = (j \operatorname{div} K \cdot K + j \bmod K) \operatorname{div} u = j \operatorname{div} K \cdot k + (j \bmod K) \operatorname{div} u$. As a consequence $\sigma_{u,M}(j) \bmod k = (j \bmod K) \operatorname{div} u < (K/2) \operatorname{div} u = k/2$.

Part 2: $\sigma_{u,M}(j+K/2) = (j+K/2) \bmod u \cdot \frac{M}{u} + (j+K/2) \operatorname{div} u = j \bmod u \cdot \frac{M}{u} + j \operatorname{div} u + k/2 = \sigma_{u,M}(j) + k/2$.

Part 3: $\sigma_{u,M}^{-1}(j) \bmod K = (j \bmod \frac{M}{u} \cdot u + j \operatorname{div} \frac{M}{u}) \bmod K = (j_1 \cdot u + j_0) \bmod K = (j_1 \operatorname{div} k \cdot K + j_1 \bmod k \cdot u + j_0) \bmod K = j_1 \bmod k \cdot u + j_0$, which gives $(\sigma_{u,M}^{-1}(j) \bmod K)/K = (j_1 \bmod k \cdot u + j_0)/K = (j_1 \bmod k + j_0/u)/k$.

□

Proof of Theorem 2.2.

Proof. Define $K = ku$. To prove the theorem, it is sufficient to prove that

$$S_{u,M} A_{K,M} S_{u,M}^{-1} \mathbf{y} = \operatorname{diag}(A_{k,n}^{0/u}, \dots, A_{k,n}^{(u-1)/u}) \mathbf{y}, \quad \text{for all } \mathbf{y}.$$

First note that the vector $A_{K,M}\mathbf{x}$ can be described by

$$\begin{cases} (A_{K,M}\mathbf{x})_j = x_j + w_K^{j \bmod K} x_{j+K/2}, \\ (A_{K,M}\mathbf{x})_{j+K/2} = x_j - w_K^{j \bmod K} x_{j+K/2}, \quad 0 \leq j \bmod K < K/2. \end{cases} \quad (\text{A.1})$$

Let $\mathbf{x} = S_{u,M}^{-1} \mathbf{y}$ and $S_{u,M}(A_{K,M}\mathbf{x}) = \mathbf{z}$, and substitute $x_j = y_{\sigma_{u,M}(j)}$ and $z_{\sigma_{u,M}(j)} = (A_{K,M}\mathbf{x})_j$ in (A.1). This gives

$$\begin{cases} z_{\sigma_{u,M}(j)} = y_{\sigma_{u,M}(j)} + w_K^{j \bmod K} y_{\sigma_{u,M}(j+K/2)}, \\ z_{\sigma_{u,M}(j+K/2)} = y_{\sigma_{u,M}(j)} - w_K^{j \bmod K} y_{\sigma_{u,M}(j+K/2)}, \quad 0 \leq j \bmod K < K/2. \end{cases} \quad (\text{A.2})$$

Defining $l = \sigma_{u,M}(j)$ and applying Lemma A.1 to j gives the following. Part 1 of Lemma A.1 says that, if $j \bmod K < K/2$ then $l \bmod k < k/2$. Furthermore, by Part 2, $\sigma_{u,M}(j+K/2) = l+k/2$. Finally, applying Part 3 to l gives $w_K^{j \bmod K} = w_K^{\sigma_{u,M}^{-1}(l) \bmod K} = w_K^{u(l' \bmod k) + s_1} = w_k^{l' \bmod k + s_1/u}$, where $l' = l \bmod \frac{M}{u}$ and $s_1 = l \operatorname{div} \frac{M}{u}$.

Substituting the above results in (A.2) gives the following description of vector $\mathbf{z} = S_{u,M} A_{K,M} S_{u,M}^{-1} \mathbf{y}$:

$$\begin{cases} z_l = y_l + w_k^{l' \bmod k + s_1/u} y_{l+k/2}, \\ z_{l+k/2} = y_l - w_k^{l' \bmod k + s_1/u} y_{l+k/2}, \quad 0 \leq l \bmod k < k/2. \end{cases} \quad (\text{A.3})$$

Writing the index $l = s_1 \cdot n + (l' \bmod k) \cdot k + l' \bmod k$, it is easy to see that $z_l = (\text{diag}(A_{k,n}^{0/u}, A_{k,n}^{1/u}, \dots, A_{k,n}^{(u-1)/u}) \cdot \mathbf{y})_l$, proving the theorem. (See Figure 2B.) \square

REFERENCES

- [1] R. C. Agarwal and J. W. Cooley. Vectorized mixed radix discrete Fourier transform algorithms. *Proceedings of the IEEE*, 75(9):1283–1292, 1987.
- [2] M. Ashworth and A. G. Lyne. A segmented FFT algorithm for vector computers. *Parallel Computing*, 6:217–224, 1988.
- [3] A. Averbuch, E. Gabber, B. Gordissky, and Y. Medan. A parallel FFT on an MIMD machine. *Parallel Computing*, 15:61–74, 1990.
- [4] R. H. Bisseling. Basic techniques for numerical linear algebra on bulk synchronous parallel computers. In L. Vulkov, J. Waśniewski, and P. Yalamov, editors, *Workshop Numerical Analysis and its Applications 1996*, volume 1196 of *Lecture Notes in Computer Science*, pages 46–57. Springer-Verlag, Berlin, 1997.
- [5] G. Bongiovanni, P. Corsini, and G. Frosini. One-dimensional and two-dimensional generalized discrete Fourier transforms. *IEEE Trans. Acoustics, Speech, and Signal Processing*, ASSP-24:97–99, 1976.
- [6] O. Bonorden, B. Juurlink, I. von Otte, and I. Rieping. The Paderborn University BSP (PUB) library – design, implementation and performance. In *13th International Parallel Processing Symposium & 10th Symposium on Parallel and Distributed Processing (IPPS/SPDP)*, San Juan, Puerto Rico, April 1999.
- [7] W. L. Briggs and V. E. Henson. *The DFT: An Owner’s Manual for the Discrete Fourier Transform*. SIAM, Philadelphia, PA, 1995.
- [8] E. Chu and A. George. FFT algorithms and their adaptation to parallel processing. *Linear Algebra and its Applications*, 284:95–124, 1998.
- [9] J. W. Cooley and J. W. Tukey. An algorithm for machine calculation of complex Fourier series. *Mathematics of Computation*, 19:297–301, 1965.
- [10] P. Corsini and G. Frosini. Properties of the multidimensional generalized discrete Fourier transform. *IEEE Transactions on Computing*, c-28(11):819–830, 1979.
- [11] A. Dubey, M. Zubair, and C. E. Grosch. A general purpose subroutine for Fast Fourier Transform on a distributed memory parallel machine. *Parallel Computing*, 20:1697–1710, 1994.

- [12] A. Gupta and V. Kumar. The scalability of FFT on parallel computers. *IEEE Transactions on Parallel and Distributed Systems*, 4(8):922–932, 1993.
- [13] J. L. Gustafson. Reevaluating Amdahl’s law. *Communications of the ACM*, 31(5):532–533, 1988.
- [14] O. Haan. A parallel one-dimensional FFT for Cray T3E. In H. Lederer and F. Hertweck, editors, *Proceedings of the Fourth European SGI/Cray MPP Workshop*, pages 188–198, IPP, Garching, Germany, September 1998.
- [15] M. Hegland. Real and complex fast Fourier transforms on the Fujitsu VPP 500. *Parallel Computing*, 22:539–553, 1996.
- [16] J. M. D. Hill, B. McColl, D. C. Stefanescu, M. W. Goudreau, K. Lang, S. B. Rao, T. Suel, T. Tsantilas, and R. H. Bisseling. BSPlib: The BSP programming library. *Parallel Computing*, 24:1947–1980, 1998.
- [17] M. A. Inda. *Constructing Parallel Algorithms for Discrete Transforms: From FFT’s to Fast Legendre Transforms*. PhD thesis, Department of Mathematics, Utrecht University, Utrecht, The Netherlands, March 2000.
- [18] M. A. Inda and R. H. Bisseling. Mixed-radix parallel fast Fourier transform. in preparation.
- [19] M. A. Inda, R. H. Bisseling, and D. K. Maslen. Parallel fast Legendre transform. In W. Zwiefelhofer and N. Kreitz, editors, *Towards Teracomputing. Proceedings of the Eighth ECMWF Workshop on the Use of Parallel Processors in Meteorology*, Reading, UK, November 1998, pages 87–108. World Scientific, Singapore, 1999.
- [20] R. Kosloff. Quantum molecular dynamics on grids. In R. E. Wyatt and J. Z. Zhang, editors, *Dynamics of Molecules and Chemical Reactions*, pages 185–230. Marcel Dekker, New York, 1996.
- [21] V. Kumar, A. Grama, A. Gupta, and G. Karypis. *Introduction to Parallel Computing: Design and Analysis of Algorithms*. The Benjamin/Cummings Publishing Company, Inc., Redwood City, CA, 1994.
- [22] W. F. McColl. Scalability, portability and predictability: The BSP approach to parallel programming. *Future Generation Computer Systems*, 12:265–272, 1996.
- [23] M. C. Pease. An adaptation of the fast Fourier transform for parallel processing. *Journal of the ACM*, 15(2):252–264, 1968.
- [24] P. A. Regalia and S. K. Mitra. Kronecker products, unary matrices and signal processing applications. *SIAM Review*, 31(4):586–613, 1989.
- [25] P. N. Swarztrauber. Multiprocessor FFTs. *Parallel Computing*, 5:197–210, 1987.
- [26] L. G. Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, 1990.
- [27] C. van Loan. *Computational Frameworks for the Fast Fourier Transform*. SIAM, Philadelphia, PA, 1992.

- [28] S. Zoldi, V. Ruban, A. Zenchuk, and S. Burtsev. Parallel implementation of the split-step Fourier method for solving nonlinear Schrödinger systems. *SIAM News*, January/February:8–9, 1999.

DEPARTMENT OF MATHEMATICS, UNIVERSITEIT UTRECHT, PO Box 80010, 3508 TA, UTRECHT,
THE NETHERLANDS

E-mail address: `inda@math.uu.nl`

DEPARTMENT OF MATHEMATICS, UNIVERSITEIT UTRECHT, PO Box 80010, 3508 TA, UTRECHT,
THE NETHERLANDS

E-mail address: `Rob.Bisseling@math.uu.nl`