

Automatic Generation of Fast Discrete Signal Transforms

Sebastian Egner, Markus Püschel

Abstract— This paper presents an algorithm that derives fast versions for a broad class of discrete signal transforms symbolically. The class includes but is not limited to the discrete Fourier and the discrete trigonometric transforms. This is achieved by finding fast sparse matrix factorizations for the matrix representations of these transforms. Unlike previous methods, the algorithm is entirely automatic and uses the defining matrix as its sole input. The sparse matrix factorization algorithm consists of two steps: first, the “symmetry” of the matrix is computed in the form of a pair of group representations; second, the representations are stepwise decomposed, giving rise to a sparse factorization of the original transform matrix. We have successfully demonstrated the method by computing automatically efficient transforms in several important cases: for the DFT, we obtain the Cooley/Tukey FFT; for a class of transforms including the DCT, type II, the number of arithmetic operations for our fast transforms is the same as for the best known algorithms. Our approach provides new insights and interpretations for the structure of the considered signal transforms and why fast algorithms do exist. The sparse matrix factorization algorithm is implemented within the software package AREP.

Keywords— discrete Fourier transform, trigonometric transforms, discrete cosine transform, fast algorithm, symmetry, group representations, monomial representations.

I. INTRODUCTION

FAST algorithms for discrete signal transforms have been a major research topic in the last decades leading to a large number of publications. Because of their wide-spread applications in digital signal processing, particular effort has been spent on the discrete Fourier transform (DFT) and the different types of trigonometric transforms, i.e., discrete cosine and sine transforms (DCTs and DSTs), as classified by Wang, [1]. Important algorithms for the DFT include the “fast Fourier transform” (FFT) found by Cooley and Tukey (first discovered by Gauß, [2]), [3], Rader’s algorithm for prime size, [4], Winograd’s algorithms, [5], as well as [6], [7], [8]. An overview on FFT algorithms can be found in [9] or [10]. Important algorithms for the trigonometric transforms were found by Chen, Smith and Fralick, [11], Wang, [12], Yip and Rao, [13], [14], Vetterli and Nussbaumer, [8], Lee, [15], Feig, [16], Chan and Ho, [17], Steidl and Tasche, [18], and Feig and Winograd, [19].

Most of the algorithms cited above are given as a factorization of the respective transform matrix into a product of

The work of Markus Püschel was supported by NSF through award 9988296 and by DARPA through research grant DABT63-98-1-0004 administered by the Army Directorate of Contracting.

S. Egner is with the Philips Research Laboratories, Eindhoven, Netherlands. E-mail: sebastian.egner@philips.com .

M. Püschel is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh. E-mail: pueschel@ece.cmu.edu .

highly structured, sparse matrices. If an algorithm is given in another way, e.g., by equations, it is possible to rewrite the algorithm in the form of a sparse matrix product.

All of these algorithms have been found by insightful manipulation of the entries of the transform matrices using algebraic relationships of these numbers. In some papers, these relationships have been referred to as “symmetry”. Several questions remain unanswered. Is there a general mathematical principle behind these algorithms, i.e., matrix factorizations? What is the appropriate definition of symmetry that accounts for the existence of the algorithms? Is it possible to automate the process of finding algorithms? For the DFT, the first two questions have been answered, as we will briefly discuss in the next subsection, since it marks the starting point for our results.

In this paper we present the mathematical background and the algorithm to automatically generate fast algorithms, given as sparse matrix factorizations, for a large class of discrete signal transforms using techniques from group representation theory. In particular we present

- An appropriate definition of “symmetry” which catches redundancy contained in the transform matrix and connects it to group representations. Furthermore, the symmetry has an intuitive interpretation in terms of signal processing. As we will see, this definition of symmetry generalizes the well-known property of the DFT diagonalizing the cyclic shift.
- An algorithm that (1) finds the symmetry of a matrix, and (2) uses it to derive a sparse matrix factorization. The algorithm has been implemented and can be used as a discover tool for fast transforms.
- The successful application of the factorization algorithm to a large class of transforms. In many cases, the generated fast transforms are similar to, or have the same arithmetic cost (operations count), as the best known algorithms.

Taken together, we provide a unifying framework that shows that a large class of the best-known fast algorithms for different transforms are all special instances of the same common principle. Thus, we shed new light on fast algorithms, put them into a common context, and give insight into their algebraic structure.

A. Signal Transforms and Group Representations

The use of finite groups and their representations is not new in signal processing. The most important example is the DFT and its connection to cyclic groups and their regular representations. This connection has been used to derive and explain the structure of the Cooley/Tukey FFT, [20], [21]. Generalization to arbitrary groups, also

known as Fourier analysis on groups, has lead to a rich class of transforms, which, however, have found no significant applications in signal processing, [22], [23], [24], [25]. An exception might be the recent paper [26], where non-regular representations of so-called wreath-product groups have been proposed for multi-resolution image processing.

The crucial step to capture in the group representation framework a broader class of signal transforms, including the cosine and sine transforms, is to leave the domain of “regular” representations in favor of the larger class of “monomial” representations. The idea has its roots in the work of Minkwitz, [27], [28], and has been further developed by the authors in [29], [30], [31], [32], [33], which forms the basis for this paper. We provide the tools to investigate a *given* transform for group representation properties, and, when appropriate, factorize the transform, thus obtaining a fast algorithm.

B. Approach

The approach for generating a fast algorithm for a given signal transform, which is given as a matrix M , consists basically of two steps. In the first step, the “symmetry” of M is computed. The “symmetry” is a pair of group representations representing an invariance property of M (cf. Section III). In the second step, the group representations are decomposed stepwise. This gives rise to factorized decomposition matrices and determines a factorization of M as a product of sparse matrices (cf. Section IV). The factorization represents a fast algorithm for the transform M . Intuitively speaking, the “symmetry” captures a large part of the redundancy contained in M , and the decomposition of the representations turns the redundancy into a fast algorithm.

C. Organization of the Paper

In Section II we introduce our notation for representing structured matrices and present the basic terms of group representations that are necessary to understand our approach for obtaining matrix factorizations. We emphasize the concepts and the methodology rather than explaining the technical details. The notion of “symmetry” of a matrix is defined in Section III and Section IV explains, how a symmetry can be used to derive a matrix factorization. In Section V we apply the matrix factorization algorithm to the Fourier transform, cosine and sine transforms of different types, the Hartley transform and the Haar transform. We compare the structure and arithmetic cost of the algorithms that we derive to the structure and cost of well-known algorithms from the literature. We conclude the paper with a brief summary and an outlook for future research in Section VI.

II. MATHEMATICAL BACKGROUND

In this section we present the basic notation of matrices and group representations we are going to use. For further information on representation theory we refer the reader to introductory books such as [34].

A. Matrices

We use the following notation to represent matrices. $[\sigma, n]$ is the $(n \times n)$ -permutation matrix corresponding to the permutation σ which is given in cycle notation, e.g., $\sigma = (1, 2, 3)(4, 6)$ means $\sigma(1) = 2, \sigma(2) = 3, \sigma(3) = 1, \sigma(4) = 6, \sigma(5) = 5, \sigma(6) = 4$ with corresponding (6×6) -matrix

$$[(1, 2, 3)(4, 6), 6] = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}.$$

Note that it is necessary to supply the size n of a permutation matrix in $[\sigma, n]$, since fixed points are omitted in cycle notation (i.e., there is a difference between $[(1, 2), 2]$ and $[(1, 2), 3]$). We prefer cycle notation because one can read off the order and the fixed points of a permutation immediately. With $\mathbf{1}_n$ we denote the identity matrix of size n , $\text{diag}(L)$ is a diagonal matrix with the list L on the diagonal. A *monomial* matrix (sometimes called scaled permutation matrix) has exactly one non-zero entry in every row and column and is represented as $[\sigma, L] = [\sigma, \text{length}(L)] \cdot \text{diag}(L)$, e.g.,

$$[(1, 2, 3), (-1, 1, 2)] = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ -1 & 0 & 0 \end{bmatrix},$$

i.e., the list L scales the columns of the matrix. The operator \otimes denotes the Kronecker (or tensor) product of matrices, and \oplus denotes the direct sum,

$$A \oplus B = \begin{bmatrix} A & \mathbf{0} \\ \mathbf{0} & B \end{bmatrix},$$

where $\mathbf{0}$ is an all-zero matrix of appropriate size.

$$R_\alpha = \begin{bmatrix} \cos(\alpha) & \sin(\alpha) \\ -\sin(\alpha) & \cos(\alpha) \end{bmatrix}$$

is the rotation matrix for angle α and

$$\text{DFT}_n = [\omega_n^{k\ell} \mid k, \ell = 0..n-1],$$

where $\omega_n = e^{2\pi j/n}$, denotes the discrete Fourier transform of size n .

B. Groups and Representations

In this paper, essentially only two types of groups will appear. The cyclic group of size n , written as $Z_n = \{1, x, \dots, x^{n-1}\}$ or, by generators and relations, as $Z_n = \langle x \mid x^n = 1 \rangle$ and the dihedral group of size $2n$ denoted by $D_{2n} = \{1, x, \dots, x^{n-1}, y, yx, \dots, yx^{n-1}\} = \langle x, y \mid x^n = y^2 = 1, y^{-1}xy = x^{-1} \rangle$.

A *representation* of a group G (over \mathbb{C}) is a homomorphism

$$\phi: G \rightarrow \text{GL}_n(\mathbb{C})$$

of G into the group $\text{GL}_n(\mathbb{C})$ of invertible $(n \times n)$ -matrices over the complex numbers \mathbb{C} . n is called the *degree* of ϕ .

Dealing with representations is nothing but dealing with groups of matrices. If ϕ is a representation of G then $\phi(G)$ is a matrix group, and, vice-versa, every matrix group can be viewed as a representation of itself. If $A \in \text{GL}_n(\mathbb{C})$, then $\phi^A : g \mapsto A^{-1} \cdot \phi(g) \cdot A$ is called the *conjugate* of ϕ by A . The representations ϕ and ϕ^A are called *equivalent*. If ϕ and ψ are representations of G , then the representation $\phi \oplus \psi : g \mapsto \phi(g) \oplus \psi(g)$ is called the *direct sum* of ϕ and ψ . The direct sum of n representations ϕ_1, \dots, ϕ_n is defined analogously. The representation ϕ is *irreducible*, if it cannot be conjugated to be a direct sum.

Theorem 1 (Maschke) *Every representation ϕ (over \mathbb{C}) of a group G can be decomposed into a direct sum of irreducible representations by conjugation with a suitable matrix A ,*

$$\phi^A = \rho_1 \oplus \dots \oplus \rho_k.$$

The ρ_i are uniquely determined up to equivalence and up to a permutation of ρ_1, \dots, ρ_k .

In other words, Theorem 1 tells us, how far a finite group of matrices can be simultaneously block diagonalized. The matrix A in Theorem 1 is not uniquely determined and is called a *decomposition matrix* for ϕ .

ϕ is called a *permutation* representation, if all images $\phi(g)$ are permutation matrices, ϕ is called a *monomial* representation, if all images $\phi(g)$ are monomial matrices. Every permutation representation is also a monomial representation.

The following example states the interpretation of the DFT in terms of representation theory.

Example 1 *It is a known fact that DFT_n maps the cyclic shift (and all its powers) in the time-domain into a phase change in the frequency-domain. In our notation,*

$$[(1, 2, \dots, n), n] \cdot \text{DFT}_n = \text{DFT}_n \cdot \text{diag}(1, \omega_n, \dots, \omega_n^{n-1}).$$

In terms of representation theory, DFT_n decomposes the permutation representation $\phi : x \mapsto [(1, 2, \dots, n), n]$ of the cyclic group $G = \mathbb{Z}_n = \langle x \mid x^n = 1 \rangle$ into the direct sum $\phi^{\text{DFT}_n} = \rho_1 \oplus \dots \oplus \rho_n$ where the irreducible representations are $\rho_k = x \mapsto \omega_n^{k-1}$.

III. SYMMETRY OF A MATRIX

The notion of symmetry has a two-fold purpose. First, it catches the redundancy contained in the matrix M ; second, it establishes the connection to representation theory, which enables the application of algebraic methods to factorize M , as sketched in Section IV.

We consider an arbitrary rectangular matrix $M \in \mathbb{C}^{m \times n}$. A *symmetry* of M is a pair (ϕ_1, ϕ_2) of representations of the same group G satisfying

$$\phi_1(g) \cdot M = M \cdot \phi_2(g), \quad \text{for all } g \in G.$$

We call G a *symmetry group* of M . We will use a shorthand notation and write $\phi_1 \xrightarrow{M} \phi_2$.

A symmetry (ϕ_1, ϕ_2) has a very natural interpretation in terms of signal processing, if M is a discrete signal transform that we multiply from the left:

For all $g \in G$, a multiplication with $\phi_2(g)$ in the time-domain corresponds to a multiplication with $\phi_1(g)$ in the frequency-domain.

With the general definition above, however, every matrix has arbitrary many symmetries. If, for example, M is an invertible $(n \times n)$ -matrix and ϕ is any representation of degree n of a group G , then M has the symmetry (ϕ, ϕ^M) . Thus, in order to catch the redundancy contained in M , we will consider several “types” of symmetry arising from restrictions on the representations ϕ_1, ϕ_2 :

1. *Mon-irred symmetry:* ϕ_1 is monomial, ϕ_2 is a direct sum of irreducible representations. If ϕ_1 is even a permutation representation then we will also speak of *perm-irred symmetry*.
2. *Mon-mon symmetry:* ϕ_1 and ϕ_2 are monomial. If ϕ_1 and ϕ_2 both are even permutation representations then we will also speak of *perm-perm symmetry*.

In words, the matrix M has a mon-mon symmetry if there are non-trivial monomial matrices L, R such that $L \cdot M = M \cdot R$. Correspondingly, the matrix M has a mon-irred symmetry if M is a decomposition matrix for a monomial representation ϕ . The rationale for considering the types of symmetry above will become clear in Section IV. Of course, one could also consider an irred-mon symmetry where ϕ_2 is monomial and ϕ_1 is decomposed. Since transposition of a matrix with irred-mon symmetry yields a matrix with mon-irred symmetry, we will restrict to the latter symmetry type. Finding symmetry of the types above is a difficult combinatorial problem and a main topic of [29], [32]. In fact, even computing the perm-perm symmetry has a complexity that is not lower than testing graphs isomorphism, which is known to be hard [35]. However, for matrices originating from signal transformations it is often practical to compute the symmetry because they contain many different entries which reduces the search space.

Example 2 *Example 1 shows that the DFT_n has the symmetry group $G = \mathbb{Z}_n = \langle x \mid x^n = 1 \rangle$ with symmetry (ϕ_1, ϕ_2) :*

$$\begin{aligned} \phi_1 : x &\mapsto [(1, 2, \dots, n), n], \\ \phi_2 : x &\mapsto \text{diag}(1, \omega_n, \dots, \omega_n^{n-1}). \end{aligned}$$

Note, that (ϕ_1, ϕ_2) is a mon-irred symmetry (even a perm-irred symmetry) as well as a mon-mon symmetry.

IV. MATRIX FACTORIZATION

Now we explain how to factorize a given matrix M , which has an arbitrary symmetry (ϕ_1, ϕ_2) . First, the representations ϕ_1, ϕ_2 are decomposed with matrices A_1, A_2 , respectively. This gives rise to two decomposed representations $\rho_1 = \phi_1^{A_1}, \rho_2 = \phi_2^{A_2}$. Second, the matrix $D = A_1^{-1} \cdot M \cdot A_2$ is computed to obtain the commutative diagram in Figure 1.




Fig. 1. Factorization of the matrix M with symmetry (ϕ_1, ϕ_2)

Altogether, we obtain the factorization

$$M = A_1 \cdot D \cdot A_2^{-1}. \quad (1)$$

From representation theory we know that D is a sparse matrix (cf. [31], Theorem 1.48, iv), but the question of sparsity remains regarding the matrices A_1 and A_2 . The factorization in (1) is useful only if the decomposition matrices A_1 and A_2 can themselves be determined as a product of sparse matrices. This is possible for monomial representations (with certain restrictions on the symmetry group G), as has been developed in the thesis research [31], [33], and justifies the consideration of the two types of symmetry described in Section III:

1. Mon-mon symmetry: A_1 and A_2 are decomposition matrices of monomial representations.
2. Mon-irred symmetry: A_1 is a decomposition matrix of a monomial representation, A_2 is the identity, since ϕ_2 is already decomposed.

In fact, we will slightly relax the definition of mon-irred symmetry and allow A_2 to be any permutation matrix, which means that ϕ_2 is a *permuted* direct sum of irreducible representations. The factorization of a decomposition matrix for a monomial representation ϕ arises from an algorithm that stepwise decomposes ϕ along a chain of normal subgroups using recursion formulas for the decomposition matrices, [33]. The recursion formula essentially determines the structure of the matrix factorizations that we will present in Section V.

The algorithm for factorizing a matrix with symmetry follows Figure 1 and reads as follows.

Algorithm 1 *Given a matrix M to be factorized into a product of sparse matrices.*

1. Determine a suitable symmetry (ϕ_1, ϕ_2) of M .
2. Decompose ϕ_1 and ϕ_2 stepwise, and obtain (factorized) decomposition matrices A_1, A_2 .
3. Compute the sparse matrix $D = A_1^{-1} \cdot M \cdot A_2$.

Result: $M = A_1 \cdot D \cdot A_2^{-1}$ is a factorization of M into a product of sparse matrices. This is a fast algorithm for evaluating the linear transformation $x \mapsto M \cdot x$.

Algorithm 1 is implemented in the library AREP, [36], a GAP share package for symbolic computation with group representations and structured matrices. GAP, [37], is a computer algebra system for symbolic computation with

groups. AREP has been created as part of the thesis research [29], [30].

In then appendix we provide an overview of Step 1 and Step 2. A comprehensive treatment including the mathematical background and all technical details can be found in [32], [33], [31], [29]. In the following we will concentrate on how the combinatorial search in Step 1 and the algebraic decomposition of Step 2 can be combined to automatically generate fast signal transforms. In particular, we are interested in answering the following questions:

- To which transforms is our approach applicable?
- What are the symmetry properties found?
- How do our generated algorithms compare to algorithms known from literature?

First, we start with two brief initial examples applying Algorithm 1 to the DFT and a circulant matrix. A more detailed version of Example 4 can be found in the appendix.

Example 3 *Let $M = \text{DFT}_4$. M has the mon-irred resp. perm-irred symmetry*

$$\begin{aligned}
\phi_1 &: x \mapsto [(1, 2, 3, 4), 4], \\
\phi_2 &: x \mapsto \text{diag}(1, \omega_4, \omega_4^2, \omega_4^3).
\end{aligned}$$

(cf. Example 2). ϕ_2 is already decomposed, hence $A_2 = \mathbf{1}_4$. Decomposing ϕ_1 stepwise yields the decomposition matrix

$$\begin{aligned}
A_1 &= \\
&(\text{DFT}_2 \otimes \mathbf{1}_2) \cdot \text{diag}(1, 1, 1, \omega_4) \cdot (\mathbf{1}_2 \otimes \text{DFT}_2) \cdot [(2, 3), 4].
\end{aligned}$$

We compute $D = A_1^{-1} \cdot M \cdot A_2 = \mathbf{1}_4$ and get the Cooley/Tukey factorization $M = A_1$.

Example 4 *Consider the circulant (4×4) -matrix*

$$M = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 4 & 1 & 2 & 3 \\ 3 & 4 & 1 & 2 \\ 2 & 3 & 4 & 1 \end{bmatrix}.$$

M has the mon-mon (even perm-perm) symmetry

$$\begin{aligned}
\phi_1 &: x \mapsto [(1, 2, 3, 4), 4], \\
\phi_2 &: x \mapsto [(1, 2, 3, 4), 4].
\end{aligned}$$

Decomposing ϕ_1 and ϕ_2 into a direct sum of irreducible representations yields the decomposition matrices

$$\begin{aligned}
A_1 = A_2 = \text{DFT}_4 = \\
(\text{DFT}_2 \otimes \mathbf{1}_2) \cdot \text{diag}(1, 1, 1, \omega_4) \cdot (\mathbf{1}_2 \otimes \text{DFT}_2) \cdot [(2, 3), 4].
\end{aligned}$$

We compute $D = A_1^{-1} \cdot M \cdot A_2 = \text{diag}(a, b, c, d)$ with complex numbers a, b, c, d (whose actual values are not important here) and obtain the well-known factorization of the cyclic convolution

$$M = \text{DFT}_4 \cdot D \cdot \text{DFT}_4^{-1}.$$

(We detail Step 1 and Step 2 in the appendix.)

V. EXAMPLES

In this section we apply Algorithm 1 to a number of signal transforms. The following factorizations have been generated from the respective transform *entirely automatically* using the GAP share package AREP, [36], [37], which contains an implementation of Algorithm 1. Even the LaTeX-expressions displayed below have been generated verbatim as they are.

We show the symmetry of the considered transforms, state the number of arithmetic operation needed by our derived fast algorithms, and compare them to algorithms known from literature. We want to emphasize that the symmetries themselves are of interest since the fast algorithms that we derive owe their existence to the symmetry in the same way as the Cooley/Tukey FFT owes its existence to the fact that the DFT diagonalizes the cyclic shift. The algebraic structure of the fast algorithms found is due to the recursion formula for decomposing monomial representations (done in Step 2 of Algorithm 1), which is subject of Theorem 2 in the Appendix.

First, we want to say some words about how AREP deals with structured matrices. AREP does *symbolic* computation with matrices which means it stores and manipulates expressions representing matrices rather than the matrices themselves. An expression is something like $(\text{DFT}_4 \otimes \mathbf{1}_{24}) \oplus \mathbf{1}_{17}$ which can be stored and manipulated more efficiently than the large matrix it represents. Of course, an expression can always be converted into a real matrix if desired. While building the structured matrices shown below (in particular in Step 2 of Algorithm 1) AREP simplifies according to specified rules. E.g., monomial or permutation matrices are extracted from Kronecker products or direct sums and multiplied if they are adjacent. In addition, occurring sparse matrices are converted into permuted direct sums or Kronecker products, if possible, to obtain a concise representation. As an example consider the sparse matrix

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 \\ 1 & 0 & 0 & 0 & 0 & -1 \end{bmatrix}.$$

M can be permuted to be a direct sum of DFT_2 's:

$$[(2, 3, 5, 4, 6), 6] \cdot (\text{DFT}_2 \oplus \text{DFT}_2 \oplus \text{DFT}_2) \cdot [(2, 6, 4, 5, 3), 6],$$

or, even more concisely,

$$[(2, 3, 5, 4, 6), 6] \cdot (\mathbf{1}_3 \otimes \text{DFT}_2) \cdot [(2, 6, 4, 5, 3), 6].$$

Note that it is an easy task to convert this expression into the original matrix above. In the same way, all permutations in the following expressions (apart from the first or the last one) can be removed leaving a product of sparse matrices each of which can be computed in-place. AREP also recognizes scrambled rotation matrices and makes this property explicit. E.g., a matrix of the form

$$R = \begin{bmatrix} -\sin(\alpha) & \cos(\alpha) \\ \cos(\alpha) & \sin(\alpha) \end{bmatrix}$$

would be transformed into the expression

$$[(1, 2), 2] \cdot R_\alpha.$$

Every matrix expression represents an algorithm for performing a matrix-vector multiplication. The number of multiplications and additions/subtractions required by this algorithm can easily be determined. E.g., the matrices

$$\text{DFT}_2, \mathbf{1}_2 \otimes \text{DFT}_2, \mathbf{1}_2 \oplus \text{DFT}_2, \frac{1}{2} \cdot \text{DFT}_2,$$

require 2 additions, 4 additions, 2 additions, 2 multiplications and 2 additions, respectively. Multiplications with -1 are not counted. Rotation matrices R_α , and scalar multiples thereof, are thought of being realized with 3 multiplications and 3 additions according to the known factorization

$$\begin{bmatrix} x & y \\ -y & x \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & -1 & 1 \end{bmatrix} \cdot \begin{bmatrix} x-y & 0 & 0 \\ 0 & y & 0 \\ 0 & 0 & x+y \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

The definitions of the transforms considered follow [38]. A matrix M representing a transform is always applied from the left, $x \mapsto M \cdot x$. The runtime for generating the algorithms, i.e., matrix factorizations, was in all cases less than 40 seconds CPU time on a Pentium II, 233 MHz with 128 MB RAM, running Linux 2.0.36.

A. DFT: Cooley/Tukey

Algorithm 1 finds the Cooley/Tukey factorization of DFT_n as illustrated in Example 3 for $n = 4$.

B. Cyclic Convolution

Algorithm 1 finds the factorization of an $(n \times n)$ circulant matrix into two DFT_n 's as illustrated for $n = 4$ in Example 4. This represents a cyclic convolution.

C. DFT: Rader

The Rader FFT, [4], computes a DFT_p of prime size p using two DFTs of size $p-1$. We apply Algorithm 1 to the case $n = 5$ and find the perm-perm symmetry

$$\begin{aligned} \phi_1 : x &\mapsto [(2, 3, 5, 4), 4], \\ \phi_2 : x &\mapsto [(2, 4, 5, 3), 4] \end{aligned}$$

with cyclic symmetry group $Z_4 = \langle x \mid x^4 = 1 \rangle$. In other words, the permutation $(2, 3, 5, 4)$ in the time-domain corresponds to the permutation $(2, 4, 5, 3)$ in the frequency-domain. The symmetry leads to the Rader factorization

$$\begin{aligned} &\text{DFT}_5 = [(4, 5), 5] \cdot \\ &(\mathbf{1}_1 \oplus ((\text{DFT}_2 \otimes \mathbf{1}_2) \cdot \text{diag}(1, 1, 1, \omega_4) \cdot (\mathbf{1}_2 \otimes \text{DFT}_2))) \cdot \\ &[(1, 4)(2, 5, 3), (a, b, c, 1, 1)] \cdot (\mathbf{1}_3 \oplus \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & -1 \end{bmatrix}) \cdot [(1, 4)(2, 3, 5), 5] \cdot \\ &(\mathbf{1}_1 \oplus \frac{1}{4} \cdot (\mathbf{1}_2 \otimes \text{DFT}_2) \cdot \text{diag}(1, 1, 1, -\omega_4) \cdot (\text{DFT}_2 \otimes \mathbf{1}_2)) \cdot \\ &[(3, 4, 5), 5]. \end{aligned}$$

The first two lines contain the matrix A_1 (essentially a DFT_4), the last two lines the matrix A_2^{-1} (essentially an inverse DFT_4), and the middle line contains the matrix D from Algorithm 1. (a, b, c) are complex constants whose actual value has been omitted for the sake of clarity.)

D. DCT, type II and III

The discrete cosine transform of type III, $\text{DCT}^{(\text{III})}$, is defined as the matrix

$$\text{DCT}_n^{(\text{III})} = \left[\sqrt{\frac{2}{n}} \cdot a_\ell \cdot \cos\left(\frac{(2k+1)\ell\pi}{2n}\right) \mid k, \ell = 0..n-1 \right],$$

where $a_k = 1/\sqrt{2}$ for $k = 0$ and $a_k = 1$ elsewhere. $\text{DCT}^{(\text{II})}$ is the transpose of $\text{DCT}^{(\text{III})}$. We compute a perm-irred symmetry for $\text{DCT}_8^{(\text{III})}$ with dihedral symmetry group $G = \text{D}_{16} = \langle x, y \mid x^8 = y^2 = 1, y^{-1}xy = x^{-1} \rangle$ and representations

$$\begin{aligned} \phi_1 : x &\mapsto [(1, 3, 5, 7, 8, 6, 4, 2), 8], & y &\mapsto [(2, 3)(4, 5)(6, 7), 8] \\ \phi_2 : x &\mapsto M_1, & y &\mapsto M_2, \end{aligned}$$

where, using $c_k = \cos(k\pi/8)$ and $s_k = \sin(k\pi/8)$,

$$M_1 = \begin{bmatrix} c_0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c_2 & 0 & 0 & 0 & 0 & 0 & s_2 \\ 0 & 0 & c_4 & 0 & 0 & 0 & s_4 & 0 \\ 0 & 0 & 0 & c_6 & 0 & s_6 & 0 & 0 \\ 0 & 0 & 0 & 0 & c_8 & 0 & 0 & 0 \\ 0 & 0 & 0 & s_{10} & 0 & c_{10} & 0 & 0 \\ 0 & 0 & s_{12} & 0 & 0 & 0 & c_{12} & 0 \\ 0 & s_{14} & 0 & 0 & 0 & 0 & 0 & c_{14} \end{bmatrix},$$

$$M_2 = \begin{bmatrix} c_0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c_1 & 0 & 0 & 0 & 0 & 0 & s_1 \\ 0 & 0 & c_2 & 0 & 0 & 0 & s_2 & 0 \\ 0 & 0 & 0 & c_3 & 0 & s_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & s_4 & 0 & 0 & 0 \\ 0 & 0 & 0 & s_5 & 0 & c_5 & 0 & 0 \\ 0 & 0 & s_6 & 0 & 0 & 0 & c_6 & 0 \\ 0 & s_7 & 0 & 0 & 0 & 0 & 0 & c_7 \end{bmatrix}.$$

Since $\text{DCT}^{(\text{III})} \cdot x$ is equivalent to $x^T \cdot \text{DCT}^{(\text{II})}$ we get the following interpretation for the $\text{DCT}^{(\text{II})}$. Permuting with $(1, 3, 5, 7, 8, 6, 4, 2)$ or $(2, 3)(4, 5)(6, 7)$ in the time-domain corresponds to multiplication with M_1 or M_2 , respectively, in the frequency-domain.

The symmetry leads to the following factorization of $\text{DCT}_8^{(\text{III})}$, which already has been derived by Minkwitz using a preliminary version of Algorithm 1 as sketched in [27], [28]. By symbolic transposition (the order of the product is reversed and each factor transposed using mathematical properties) we get a factorization of $\text{DCT}_8^{(\text{II})}$ which essentially corresponds to the first algorithm, found by Chen, Smith and Fralick, [11], which directly computes the $\text{DCT}^{(\text{II})}$ without using the DFT.

$$\begin{aligned} \text{DCT}_8^{(\text{III})} &= \\ &[(1, 2, 6, 8)(3, 7, 5, 4), 8] \cdot \\ &(\mathbf{1}_2 \otimes ((\mathbf{1}_2 \otimes \text{DFT}_2) \cdot [(2, 3), 4] \cdot (\text{DFT}_2 \oplus \mathbf{1}_2))) \cdot \\ &[(2, 7, 6, 8, 5, 4, 3), 8] \cdot (\mathbf{1}_4 \oplus \frac{1}{\sqrt{2}} \cdot \text{DFT}_2 \oplus \mathbf{1}_2) \cdot \\ &[(5, 6), 8] \cdot ((\text{DFT}_2 \otimes \mathbf{1}_3) \oplus \mathbf{1}_2) \cdot [(2, 8, 3, 7, 4), 8] \cdot \\ &\frac{1}{2} \cdot (\frac{1}{\sqrt{2}} \cdot \mathbf{1}_2 \oplus \text{R}_{\frac{3}{8}\pi} \oplus \text{R}_{\frac{15}{16}\pi} \oplus \text{R}_{\frac{11}{16}\pi}) \cdot \\ &[(2, 5)(4, 7)(6, 8), 8], \end{aligned}$$

(13 mults, 29 adds),

and, by transposition,

$$\begin{aligned} \text{DCT}_8^{(\text{II})} &= \\ &[(2, 5)(4, 7)(6, 8), 8] \cdot \\ &\frac{1}{2} \cdot (\frac{1}{\sqrt{2}} \cdot \mathbf{1}_2 \oplus \text{R}_{\frac{3}{8}\pi} \oplus \text{R}_{\frac{15}{16}\pi} \oplus \text{R}_{\frac{21}{16}\pi}) \cdot \\ &[(2, 4, 7, 3, 8), 8] \cdot ((\text{DFT}_2 \otimes \mathbf{1}_3) \oplus \mathbf{1}_2) \cdot [(5, 6), 8] \cdot \\ &(\mathbf{1}_4 \oplus \frac{1}{\sqrt{2}} \cdot \text{DFT}_2 \oplus \mathbf{1}_2) \cdot [(2, 3, 4, 5, 8, 6, 7), 8] \cdot \\ &(\mathbf{1}_2 \otimes ((\text{DFT}_2 \oplus \mathbf{1}_2) \cdot [(2, 3), 4] \cdot (\mathbf{1}_2 \otimes \text{DFT}_2))) \cdot \\ &[(1, 8, 6, 2)(3, 4, 5, 7), 8], \end{aligned}$$

(13 mults, 29 adds).

Looking at the factorization of $\text{DCT}_8^{(\text{III})}$, the first four lines give the matrix A_1 from Algorithm 1, the last line contains the permutation matrix A_2^{-1} (which makes the block structure of M_1 and M_2 explicit), and the fifth line gives the matrix D .

The algorithms for $\text{DCT}_8^{(\text{III})}$ and $\text{DCT}_8^{(\text{II})}$ have the same arithmetic cost as the best known algorithms, [11], [12], [8], [15], [39], [17], [18], [19]. Note that those who use only 12 multiplications do not normalize the first row of the $\text{DCT}_8^{(\text{II})}$, which saves one multiplication. The only algorithm that claims 11 multiplications, [40], considers a scaled version of the $\text{DCT}^{(\text{II})}$ matrix,

$$\text{DCT}^{(\text{II})}'_n = \sqrt{2} \cdot \sqrt{\frac{n}{2}} \cdot \text{DCT}^{(\text{II})}_n.$$

Multiplying by scalars conserves the perm-irred symmetry (it just changes the matrix D in Algorithm 1) and AREP also finds a factorization with 11 multiplications,

$$\begin{aligned} \text{DCT}^{(\text{II})}'_8 &= \\ &[(2, 5)(4, 7)(6, 8), 8] \cdot \\ &(\mathbf{1}_2 \oplus \sqrt{2} \cdot \text{R}_{\frac{3}{8}\pi} \oplus \sqrt{2} \cdot \text{R}_{\frac{15}{16}\pi} \oplus \sqrt{2} \cdot \text{R}_{\frac{21}{16}\pi}) \cdot \\ &[(2, 4, 7, 3, 8), 8] \cdot ((\text{DFT}_2 \otimes \mathbf{1}_3) \oplus \mathbf{1}_2) \cdot [(5, 6), 8] \cdot \\ &(\mathbf{1}_4 \oplus \frac{1}{\sqrt{2}} \cdot \text{DFT}_2 \oplus \mathbf{1}_2) \cdot [(2, 3, 4, 5, 8, 6, 7), 8] \cdot \\ &(\mathbf{1}_2 \otimes ((\text{DFT}_2 \oplus \mathbf{1}_2) \cdot [(2, 3), 4] \cdot (\mathbf{1}_2 \otimes \text{DFT}_2))) \cdot \\ &[(1, 8, 6, 2)(3, 4, 5, 7), 8], \end{aligned}$$

(11 mults, 29 adds).

We want to mention that the $\text{DCT}^{(\text{III})}$ (and hence the $\text{DCT}^{(\text{II})}$) also has a mon-mon symmetry. E.g., for the case $n = 8$ the symmetry group is the direct product $\mathbf{Z}_2 \times \mathbf{Z}_8$. In fact, this symmetry has been used by Feig and Winograd to derive a fast $\text{DCT}^{(\text{II})}$ algorithm, [19], and a lower bound for the number of non-rational multiplications necessary for computing the $\text{DCT}^{(\text{II})}$ (for $n = 8$ the optimal number is 11, where the first row of the $\text{DCT}^{(\text{II})}$ is unscaled). They essentially follow Algorithm 1 with the difference that ϕ_1 and ϕ_2 are only decomposed over the rational numbers \mathbb{Q} (which yields a coarser decomposition of ϕ_1 and ϕ_2) using rational matrices A_1 and A_2 . All non-rational multiplications then are concentrated in the block diagonal matrix D . AREP currently is only capable to decompose representations over \mathbb{C} .

E. DCT, type IV

The discrete cosine transform of type IV, $\text{DCT}^{(\text{IV})}$, is defined as the matrix

$$\text{DCT}_n^{(\text{IV})} = \left[\sqrt{\frac{2}{n}} \cdot \cos\left(\frac{(2k+1)(2\ell+1)\pi}{4n}\right) \mid k, \ell = 0..n-1 \right].$$

We compute a mon-irred symmetry for $\text{DCT}_8^{(\text{IV})}$ with dihedral symmetry group $G = \text{D}_{32} = \langle x, y \mid x^{16} = y^2 = 1, y^{-1}xy = x^{-1} \rangle$ and representations

$$\begin{aligned} \phi_1 : x &\mapsto [(1, 3, 5, 7, 8, 6, 4, 2), (1, 1, 1, 1, 1, 1, 1, -1)], \\ &y \mapsto [(2, 3)(4, 5)(6, 7), (1, 1, 1, 1, 1, 1, 1, -1)], \\ \phi_2 : x &\mapsto M_1, \quad y \mapsto M_2, \end{aligned}$$

i.e., compared to the perm-irred symmetry of the $\text{DCT}_8^{(\text{III})}$, the last column of the images of $\phi_1(x)$ and $\phi_1(y)$ are multiplied by -1 which also leads to a larger group. Using $c_k = \cos(k\pi/16)$, $s_k = \sin(k\pi/16)$, the matrices M_1, M_2 are given by

$$M_1 = \begin{bmatrix} c_2 & 0 & 0 & 0 & 0 & 0 & 0 & s_2 \\ 0 & c_6 & 0 & 0 & 0 & 0 & s_6 & 0 \\ 0 & 0 & c_{10} & 0 & 0 & s_{10} & 0 & 0 \\ 0 & 0 & 0 & c_{14} & s_{14} & 0 & 0 & 0 \\ 0 & 0 & 0 & s_{18} & c_{18} & 0 & 0 & 0 \\ 0 & 0 & s_{22} & 0 & 0 & c_{22} & 0 & 0 \\ 0 & s_{26} & 0 & 0 & 0 & 0 & c_{26} & 0 \\ s_{30} & 0 & 0 & 0 & 0 & 0 & 0 & c_{30} \end{bmatrix},$$

$$M_2 = \begin{bmatrix} c_1 & 0 & 0 & 0 & 0 & 0 & 0 & s_1 \\ 0 & c_3 & 0 & 0 & 0 & 0 & s_3 & 0 \\ 0 & 0 & c_5 & 0 & 0 & s_5 & 0 & 0 \\ 0 & 0 & 0 & c_7 & s_7 & 0 & 0 & 0 \\ 0 & 0 & 0 & s_9 & c_9 & 0 & 0 & 0 \\ 0 & 0 & s_{11} & 0 & 0 & c_{11} & 0 & 0 \\ 0 & s_{13} & 0 & 0 & 0 & 0 & c_{13} & 0 \\ s_{15} & 0 & 0 & 0 & 0 & 0 & 0 & c_{15} \end{bmatrix}.$$

The symmetry leads to the following factorization of $\text{DCT}_8^{(\text{IV})}$. Since $\text{DCT}^{(\text{IV})}$ is symmetric, transposition leads to another factorization of $\text{DCT}_8^{(\text{IV})}$ which is very close to the fast algorithm given by Wang, [12].

$$\begin{aligned} \text{DCT}_8^{(\text{IV})} &= \\ &[(1, 2, 8)(3, 6, 5), (1, -1, 1, 1, 1, -1, 1, 1)] \cdot \\ &(\mathbf{1}_2 \otimes ((\mathbf{1}_2 \oplus \frac{1}{\sqrt{2}} \cdot \text{DFT}_2) \cdot [(3, 4), 4] \cdot (\text{DFT}_2 \otimes \mathbf{1}_2))) \cdot \\ &[(1, 3)(2, 4)(5, 7)(6, 8), 8] \cdot (\mathbf{1}_4 \oplus \text{R}_{\frac{15}{8}\pi} \oplus \text{R}_{\frac{11}{8}\pi}) \cdot \\ &(\text{DFT}_2 \otimes \mathbf{1}_4) \cdot [(3, 5, 7)(4, 6, 8), 8] \cdot \\ &\frac{1}{2} \cdot (\text{R}_{\frac{31}{32}\pi} \oplus \text{R}_{\frac{19}{32}\pi} \oplus \text{R}_{\frac{27}{32}\pi} \oplus \text{R}_{\frac{23}{32}\pi}) \cdot \\ &[(1, 8, 5, 6, 3, 2)(4, 7), 8], \end{aligned}$$

(20 mults, 38 adds).

The first four lines correspond to the decomposition matrix A_1 of ϕ_1 in Algorithm 1, the fifth line to D , and the last line contains the permutation matrix A_2^{-1} , the inverse of which permutes ϕ_2 to be a direct sum.

An algorithm with two additions less can be found in [17].

F. DST, type II and III

The discrete sine transform of type III, $\text{DST}^{(\text{III})}$, is defined as the matrix

$$\text{DST}_n^{(\text{III})} = \left[\sqrt{\frac{2}{n}} \cdot a_k \cdot \sin\left(\frac{(2k+1)(\ell+1)\pi}{2n}\right) \mid k, \ell = 0..n-1 \right],$$

where $a_k = 1/\sqrt{2}$ for $k = 0$ and $a_k = 1$ elsewhere. $\text{DST}^{(\text{II})}$ is the transpose of $\text{DST}^{(\text{III})}$. We compute a perm-irred symmetry for $\text{DST}_8^{(\text{III})}$ with dihedral symmetry group $G = \text{D}_{16} = \langle x, y \mid x^8 = y^2 = 1, y^{-1}xy = x^{-1} \rangle$ and representations

$$\begin{aligned} \phi_1 : x &\mapsto [(1, 3, 5, 7, 8, 6, 4, 2), (-1, 1, 1, 1, 1, 1, 1, -1)], \\ &y \mapsto [(2, 3)(4, 5)(6, 7), (-1, 1, 1, 1, 1, 1, 1, -1)], \\ \phi_2 : x &\mapsto M_1, \quad y \mapsto M_2, \end{aligned}$$

i.e., compared to the perm-irred symmetry of the $\text{DCT}_8^{(\text{III})}$, the first and last column of the images of $\phi_1(x)$ and $\phi_1(y)$ are multiplied by -1. The matrices M_1 and M_2 (not given due to lack of space) have entries $\neq 0$ only at positions (i, j) with $i + j = 8$, $i, j = 1..8$.

The symmetry leads to a factorization of $\text{DST}_8^{(\text{III})}$, and hence to a factorization of $\text{DST}_8^{(\text{II})}$, which requires 13 additions and 29 multiplications.

Since $\text{DCT}^{(\text{II})}$, $\text{DCT}^{(\text{III})}$, $\text{DST}^{(\text{II})}$ and $\text{DST}^{(\text{III})}$ all have the same arithmetic cost (because type II and III are transposed and [17], Section 4.2), the algorithms found by AREP are among the best known algorithms.

G. DST, type IV

The discrete sine transform of type IV, $\text{DST}^{(\text{IV})}$, is defined as the matrix

$$\text{DST}_n^{(\text{IV})} = \left[\sqrt{\frac{2}{n}} \cdot \sin\left(\frac{(2k+1)(2\ell+1)\pi}{4n}\right) \mid k, \ell = 0..n-1 \right].$$

We compute a mon-irred symmetry for $\text{DCT}_8^{(\text{IV})}$ with dihedral symmetry group $G = \text{D}_{32} = \langle x, y \mid x^{16} = y^2 = 1, y^{-1}xy = x^{-1} \rangle$ and representations

$$\begin{aligned} \phi_1 : x &\mapsto [(1, 3, 5, 7, 8, 6, 4, 2), (-1, 1, 1, 1, 1, 1, 1, 1)], \\ &y \mapsto [(2, 3)(4, 5)(6, 7), (-1, 1, 1, 1, 1, 1, 1, 1)], \\ \phi_2 : x &\mapsto M_1, \quad y \mapsto M_2, \end{aligned}$$

i.e., the difference to the perm-irred symmetry of the $\text{DCT}_8^{(\text{III})}$ lies only in the first column of the images being multiplied by -1. The matrices M_1 and M_2 (not given due to lack of space) have entries $\neq 0$ only on the diagonal and opposite diagonal.

The symmetry leads to a factorization of $\text{DST}_8^{(\text{IV})}$, which requires 20 multiplications and 38 additions. As for the $\text{DCT}^{(\text{IV})}$, this is two additions more as in the best known algorithm, [17].

H. DCT and DST, type I

Although the transforms $\text{DCT}^{(\text{I})}$ and $\text{DST}^{(\text{I})}$ do not have a mon-irred symmetry, they do possess a mon-mon symmetry which can be used for their factorization. However, the algorithms obtained this way are not as good as those from [17].

I. Hartley transform

The discrete Hartley transform DHT_n is defined as the matrix

$$\text{DHT}_n = \left[\cos\left(\frac{2k\ell\pi}{n}\right) + \sin\left(\frac{2k\ell\pi}{n}\right) \mid k, \ell = 0..n-1 \right].$$

Note that we omitted the normalization factor $1/\sqrt{n}$ in the definition to obtain a fair comparison to known algorithms. The DHT_8 has a perm-irred symmetry with dihedral symmetry group $G = \text{D}_{16} = \langle x, y \mid x^8 = y^2 = 1, y^{-1}xy = x^{-1} \rangle$ and representations

$$\begin{aligned} \phi_1 : x &\mapsto [(1, 2, 3, 4, 5, 6, 7, 8), 8], & y &\mapsto [(2, 8)(3, 7)(4, 6), 8] \\ \phi_2 : x &\mapsto M_1, & y &\mapsto [(2, 8)(3, 7)(4, 6), 8]. \end{aligned}$$

The only nonzero entries of matrix M_1 are at (i, j) where $i + j = 9$. The symmetry yields the following factorization.

$$\begin{aligned} \text{DHT}_8 = & \\ & [(1, 8)(2, 4)(3, 5, 7, 6), 8] \cdot \\ & (\mathbf{1}_2 \otimes ((\mathbf{1}_2 \otimes \text{DFT}_2) \cdot [(2, 3), 4] \cdot (\text{DFT}_2 \oplus \mathbf{1}_2))) \cdot \\ & [(2, 7, 6, 8, 5, 4, 3), 8] \cdot (\mathbf{1}_4 \oplus -\frac{1}{\sqrt{2}} \cdot \text{DFT}_2 \oplus \mathbf{1}_2) \cdot \\ & ((\text{DFT}_2 \otimes \mathbf{1}_3) \oplus \mathbf{1}_2) \cdot \\ & [(2, 5, 3, 6, 4)(7, 8), (1, -1, -\sqrt{2}, -\sqrt{2}, \sqrt{2}, \sqrt{2}, 1, 1)] \cdot \\ & (\mathbf{1}_6 \oplus -\text{DFT}_2) \cdot [(2, 5, 8, 7, 3, 4), 8], \\ & \quad (6 \text{ mults, } 22 \text{ adds}). \end{aligned}$$

Closer investigation shows that two of the $\sqrt{2}$ could be canceled against the $1/\sqrt{2}$ yielding 4 multiplications less.

The best algorithm for the DHT_8 with respect to arithmetic operations seems to be the split-radix algorithm given in Sorensen et.al., [41], and also needs 2 multiplications and 22 additions.

J. Haar transform

The Haar transform HT_{2^k} is defined recursively by:

$$\text{HT}_2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad \text{HT}_{2^{k+1}} = \begin{bmatrix} \text{HT}_{2^k} & \otimes \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \\ 2^{k/2} \cdot \mathbf{1}_{2^k} & \otimes \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \end{bmatrix},$$

for $k \geq 1$. The transpose of the Haar transform has a perm-irred symmetry. The symmetry group is an iterated wreath product, [26]. For $k = 3$ we obtain the following factorization.

$$\begin{aligned} \text{HT}_8 = & \\ & [(3, 4)(5, 7)(6, 8), (\frac{1}{8}, -\frac{1}{8}, \frac{1}{4\sqrt{2}}, \frac{1}{4\sqrt{2}}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})] \cdot \\ & (\text{DFT}_2 \oplus \mathbf{1}_6) \cdot [(2, 5, 3)(4, 6), 8] \cdot \\ & (\mathbf{1}_2 \otimes ((\text{DFT}_2 \oplus \mathbf{1}_2) \cdot [(2, 3), 4] \cdot (\mathbf{1}_2 \otimes \text{DFT}_2))) \cdot \\ & [(1, 5)(2, 6)(3, 7)(4, 8), 8], \\ & \quad (8 \text{ mults, } 14 \text{ adds}). \end{aligned}$$

The first line contains the matrix D from Algorithm 1 and the other lines the matrix A_1 (we decomposed the transpose of $\text{HT}(8)$ and transposed the result). The number of operations coincides with the best known algorithm, [42].

K. Wreath Product Transform

In a recent paper, [26], decomposition matrices for permutation representations of certain groups, which are called iterated wreath products, have been proposed for image processing. By construction, these transforms possess symmetry in our definition. Thus, we readily obtain the following factorization of the (16×16) transform W_{16} given in [26], p. 117 (where it is called A ; we want to note that this transform has an error in column 15, where the last $-i$ should read i). For brevity, we denote $j = \omega_4$.

$$\begin{aligned} W_{16} = & \\ & [(1, 14, 7)(2, 15, 8)(3, 16, 5)(4, 13, 6)(9, 11)(10, 12), 16] \cdot \\ & (\mathbf{1}_2 \otimes ((\mathbf{1}_2 \otimes ((\text{DFT}_2 \otimes \mathbf{1}_2) \cdot \text{diag}(1, 1, 1, j) \cdot (\mathbf{1}_2 \otimes \text{DFT}_2))) \cdot \\ & [(2, 5)(4, 7), 8] \cdot (\text{DFT}_2 \oplus \mathbf{1}_6))) \cdot \\ & [(3, 13, 11, 15, 7, 5, 9)(4, 14, 12, 16, 8, 6, 10), \\ & (1, 1, 1, j, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1)] \cdot \\ & ((\text{DFT}_2 \otimes \mathbf{1}_2) \oplus \mathbf{1}_{12}) \cdot [(5, 14, 8, 13, 6, 16)(7, 15)(9, 10, 12), \\ & (1, j, -1, -j, -j, -1, -1, -j, -1, 1, 1, -1, j, -1, -1, j)] \\ & \quad (11 \text{ mults by } j, 24 \text{ adds}). \end{aligned}$$

VI. CONCLUSION AND FUTURE RESEARCH

We have presented an entirely automatic method for symbolically deriving fast algorithms for an important class of discrete linear signal transforms. This class includes the DFT, DCT, DST, Hartley, and Haar transform. In most cases, the derived algorithms were among the best ones known. The approach is based on the definition of ‘‘symmetry’’ of a transform as a pair of group representations, which operate in the time-domain and the frequency-domain, respectively, leaving the transform invariant. More precisely, the considered transform matrices can be factorized because they are decomposition matrices of monomial group representations.

The results of this paper open at least the following two research questions:

1. How do symmetry and signal processing properties of a transform relate to each other?
2. Is it possible to extend the approach described to derive a larger class of fast signal transforms?

AREP includes an interface to SPL. SPL is a domain specific language and compiler for (1) representing a fast signal transform given as a matrix expression like the ones generated by AREP, and (2) translate it into an efficient, e.g., FORTRAN program for computing the transform. SPL is under development within [43]. The interface between AREP and SPL allows the automatic implementation of all algorithms derived by AREP.

VII. ACKNOWLEDGMENTS

Part of the presented research has been developed when the authors were with Prof. Beth at the Institute of Algorithms and Cognitive Systems at the University of Karlsruhe, Germany. The authors are indebted to Prof. Beth

for his support that made this work possible. The authors also want to thank Prof. Moura for helpful suggestions and proof-reading.

REFERENCES

- [1] Z. Wang and B.R. Hunt, "The Discrete W Transform," *Applied Mathematics and Computation*, vol. 16, pp. 19–48, 1985.
- [2] M.T. Heideman, D.H. Johnson, and C.S. Burrus, "Gauss and the History of the Fast Fourier Transform," *Archive for History of Exact Sciences*, vol. 34, pp. 265–277, 1985.
- [3] J. W. Cooley and J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," *Math. of Computation*, vol. 19, pp. 297–301, 1965.
- [4] C. M. Rader, "Discrete Fourier Transforms When the Number of Data Samples is Prime," *Proceedings of the IEEE*, vol. 56, pp. 1107–1108, 1968.
- [5] S. Winograd, *Arithmetic Complexity of Computation*, Siam, 1980.
- [6] R. C. Agarwal and J. W. Cooley, "New Algorithms for Digital Convolution," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. ASSP-25, pp. 392–410, 1977.
- [7] L. I. Bluestein, "A Linear Filtering Approach to the Computation of the Discrete Fourier Transform," *IEEE Trans. on Audio and Electroacoustics*, vol. AU-18, no. 4, pp. 451–455, 1970.
- [8] M. Vetterli and H.J. Nussbaumer, "Simple FFT and DCT Algorithms with reduced Number of Operations," *Signal Processing*, vol. 6, pp. 267–278, 1984.
- [9] H. J. Nussbaumer, *Fast Fourier Transformation and Convolution Algorithms*, Springer, 2nd edition, 1982.
- [10] R. Tolimieri, M. An, and C. Lu, *Algorithms for Discrete Fourier Transforms and Convolution*, Springer, 2nd edition, 1997.
- [11] W.-H. Chen, C.H. Smith, and S.C. Fralick, "A Fast Computational Algorithm for the Discrete Cosine Transform," *IEEE Trans. on Communications*, vol. COM-25, no. 9, pp. 1004–1009, 1977.
- [12] Z. Wang, "Fast Algorithms for the Discrete W Transform and for the Discrete Fourier Transform," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, vol. ASSP-32, no. 4, pp. 803–816, 1984.
- [13] P. Yip and K.R. Rao, "Fast Decimation-In-Time Algorithms for a Family of Discrete Sine and Cosine Transforms," *Circuits, Systems, and Signal Processing*, vol. 3, no. 4, pp. 387–408, 1984.
- [14] P. Yip and K.R. Rao, "The Decimation-In-Frequency Algorithms for a Family of Discrete Sine and Cosine Transforms," *Circuits, Systems, and Signal Processing*, vol. 7, no. 1, pp. 3–19, 1988.
- [15] B. Lee, "A New Algorithm to Compute the Discrete Cosine Transform," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, vol. ASSP-32, no. 6, pp. 1243–1245, 1984.
- [16] E. Feig, "A fast scaled-DCT algorithm," in *SPIE Proceedings*, 1990, vol. 1244, pp. 2–13.
- [17] S.C. Chan and K.L. Ho, "Direct Methods for computing discrete sinusoidal transforms," *IEE Proceedings*, vol. 137, no. 6, pp. 433–442, 1990.
- [18] G. Steidl and M. Tasche, "A Polynomial Approach to Fast Algorithms for Discrete Fourier-Cosine and Fourier-Sine Transforms," *Mathematics of Computation*, vol. 56, no. 193, pp. 281–296, 1991.
- [19] E. Feig and S. Winograd, "Fast Algorithms for the Discrete Cosine Transform," *IEEE Trans. on Signal Processing*, vol. 40, no. 9, pp. 2174–2193, 1992.
- [20] L. Auslander, E. Feig, and S. Winograd, "Abelian Semi-simple Algebras and Algorithms for the Discrete Fourier Transform," *Advances in Applied Mathematics*, vol. 5, pp. 31–55, 1984.
- [21] Th. Beth, *Verfahren der Schnellen Fouriertransformation*, Teubner, 1984.
- [22] D. Eberly and D. Wenzel, "Adaptation of Group Algebras to Signal and Image Processing," *CVGIP: Graphical Models and Image Processing*, vol. 53, no. 4, pp. 340–348, 1991.
- [23] R.J. Valenza, "A Representation-Theoretic Approach to the DFT with Noncommutative Generalizations," *IEEE Trans. on Signal Processing*, vol. 40, no. 4, pp. 814–822, 1992.
- [24] D. Rockmore, "Some applications of generalized FFT's," in *Proceedings of DIMACS Workshop in Groups and Computation*, 1995, vol. 28, pp. 329–370.
- [25] A. Terras, *Fourier Analysis on Finite Groups and Applications*, Cambridge University Press, 1999.
- [26] R. Foote, G. Mirchandi, D. Rockmore, D. Healy, and T. Olson, "A Wreath Product Approach to Signal and Image Processing: Part I — Multiresolution Analysis," *IEEE Trans. on Signal Processing*, vol. 48, no. 1, pp. 102–132, 2000.
- [27] T. Minkwitz, *Algorithmensynthese für lineare Systeme mit Symmetrie*, Ph.D. thesis, Universität Karlsruhe, Informatik, 1993.
- [28] T. Minkwitz, "Algorithms Explained by Symmetry," *Lecture Notes on Computer Science*, vol. 900, pp. 157–167, 1995.
- [29] S. Egner, *Zur Algorithmischen Zerlegungstheorie Linearer Transformationen mit Symmetrie*, Ph.D. thesis, Universität Karlsruhe, Informatik, 1997.
- [30] M. Püschel, *Konstruktive Darstellungstheorie und Algorithmen-generierung*, Ph.D. thesis, Universität Karlsruhe, Informatik, 1998, Translated in [31].
- [31] M. Püschel, "Constructive Representation Theory and Fast Signal Transforms," Tech. Rep. Drexel-MCS-1999-1, Drexel University, Philadelphia, 1999, Translation of [30].
- [32] S. Egner and M. Püschel, "Fast Discrete Signal Transforms and Monomial Representations of Finite Groups," submitted for publication.
- [33] M. Püschel, "Decomposing Monomial Representations of Solvable Groups," Tech. Rep. Drexel-MCS-1999-2, Drexel Univ., Philadelphia, 1999, submitted for publication.
- [34] G. James and M. Liebeck, *Representations and Characters of Groups*, Cambridge Univ. Pr., 1993.
- [35] M. R. Garey and D.S. Johnson, *Computers and Intractability*, W. H. Freeman and Co., 1979.
- [36] S. Egner and M. Püschel, *AREP – Constructive Representation Theory and Fast Signal Transforms*, GAP share package, 1998, <http://avalon.ira.uka.de/home/pueschel/arep/arep.html>.
- [37] The GAP Team, University of St. Andrews, Scotland, *GAP – Groups, Algorithms, and Programming*, 1997, <http://www-gap.dcs.st-and.ac.uk/~gap/>.
- [38] K. R. Rao and P. Yip, *Discrete Cosine Transform: Algorithms, Advantages, Applications*, Academic Press, 1990.
- [39] H.S. Hou, "A Fast Recursive Algorithm For Computing the Discrete Cosine Transform," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, vol. ASSP-35, no. 10, pp. 1455–1461, 1987.
- [40] Loeffler, "Practical Fast 1-D DCT Algorithms with 11 Multiplications," in *Proceedings ICASSP*, 1989, pp. 988–991.
- [41] H.V. Sorensen, D.L. Jones, C.S. Burrus, and M.T. Heideman, "On Computing the Discrete Hartley Transform," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, vol. ASSP-33, no. 4, pp. 1231–1238, 1985.
- [42] D. F. Elliott and K. R. Rao, *Fast Transforms: Algorithms, Analyses, Applications*, Academic Press, 1982.
- [43] J. M. F. Moura, J. Johnson, R. W. Johnson, D. Padua, V. Prasanna, M. Püschel, and M. M. Veloso, "SPIRAL: Portable Library of Optimized Signal Processing Algorithms," 1998, <http://www.ece.cmu.edu/~spiral/>.
- [44] J.S. Leon, "Permutation Group Algorithms Based on Partitions, I: Theory and Algorithms," *Journal of Symbolic Computation*, vol. 12, no. 4/5, pp. 533–583, October/November 1991.

APPENDIX

In the following we will give an overview on Step 1 and 2 of Algorithm 1. For a comprehensive treatment including all technical details we have to refer the reader to [32], [33], [31], [30], [29], [27], which also provide the necessary prerequisites from representation theory of finite groups. Of particular interest might be the recursion formula for Step 2, given in Theorem 2, which gives the explanation for the algebraic structure of the fast transforms found in Section V.

A. Algorithm 1, Step 1

In Step 1 of Algorithm 1 the mon-mon or mon-irred symmetry of the matrix M has to be computed. The two types require different approaches.

For the mon-mon symmetry we are interested in the group G of all pairs (L, R) of monomial matrices such that

$LM = MR$. (The representations ϕ_1, ϕ_2 of Section III are the projections $(L, R) \mapsto L$ and $(L, R) \mapsto R$.) However, this general definition may lead to an infinite G . Therefore we artificially restrict L and R to contain k -th roots of unity: Let $\text{MonMon}_k(M)$ denote the group of all pairs of monomial matrices (L, R) such that $LM = MR$ and all entries in L and R are k -th roots of unity.

To construct $\text{MonMon}_k(M)$ we replace each entry M_{ij} by the $(k \times k)$ -matrix $M' = [M_{ij}\omega_k^{r+s} \mid r, s]$ where ω_k denotes a fixed primitive k -th root of unity. This encoding of M turns monomial operations on M (with k -th roots of unity) into permutations operations on the larger matrix M' . The parameter k is chosen dependent on the entries of M . The encoding has first been described in [44] for the special case of finite fields. It reduces MonMon_k to the well known problem of constructing all pairs (L, R) of permutations such that $LM' = M'R$ for a given matrix M' (the perm-perm symmetry of M'). The most efficient methods to solve the latter problem are partition-based backtracking methods that systematically try all permutations, removing entire branches quickly. This is described by J. Leon in [44] who also distributes a very efficient C-program to compute the symmetry. Moreover, [29] describes a program in GAP to compute MonMon_k and it is proven in [31], [30] that MonMon_k can indeed be computed the way we have claimed.

Now we turn to the other important type of symmetry, the mon-irred symmetry. Here we are interested in all monomial matrices L such that $R = MLM^{-1}$ is permuted block-diagonal for invertible M . Formally, “permuted block-diagonal” means that there is a permutation π such that $R = \pi^{-1} \text{diag}(R_1, \dots, R_m)\pi$ for smaller matrices R_1, \dots, R_m . As a quantitative measure for block-diagonality we define the *conjugated block structure* (cbs) of a matrix A as the partition $\text{cbs}(A) = \{1, \dots, n\} / \sim^*$ where \sim^* is the reflexive-symmetric-transitive closure of the relation \sim defined by $i \sim j \Leftrightarrow A_{ij} \neq 0$. The partitions are partially ordered by the refinement.

There are two approaches to the mon-irred symmetry. The first one essentially enumerates all monomial matrices L and collects them into groups according to the block structures $\text{cbs}(MLM^{-1})$ they give rise to. The result is a list of pairs (G, p) such that G is the group of all monomial matrices L such that $\text{cbs}(MLM^{-1}) \leq p$ and p is the join of all $\text{cbs}(MLM^{-1})$. Each of the groups G qualifies as a mon-irred symmetry, it only has to be tested if the matrices $R = MLM^{-1}$ form a representation of G that is a direct sum of irreducible representations.

The second approach to the mon-irred symmetry essentially enumerates all partitions p and constructs the group G of all matrices L such that $\text{cbs}(MLM^{-1})$ is a refinement of p . The result is the same list of pairs (G, p) as before. The main difference to the first method is that one can restrict the enumeration of partitions to the cases that will be most useful for the decomposition of the signal transform, namely those with many small blocks. This is much faster than running through all monomial matrices. For the details on the methods and on a number of important im-

provements see [29], [32]. All of the mon-irred symmetries used for the examples in this article are found in a few seconds on a standard workstation using our implementation in GAP for the library AREP.

B. Algorithm 1, Step 2

Step 2 of Algorithm 1 requires to decompose a given monomial representation ϕ of a group G into a direct sum of irreducible representations. In addition, the corresponding decomposition matrix has to be computed as a product of structured sparse matrices.

First we need to introduce the notion of a *transitive* monomial representation. Let $n = \deg \phi$ and denote by e_i the i -th canonical base vector (i -th entry = 1, = 0 else). Then ϕ is called transitive, if for all i, j there is a $g \in G$ such that (the monomial matrix) $\phi(g)$ maps e_i to a multiple of e_j .

The key construct for the decomposition algorithm is the *induction* of representations. In short, induction constructs a representation of G from a representation of a subgroup of G . More precisely, let $H \leq G$ be a subgroup with representation λ and let $T = (t_1, \dots, t_k)$ be a transversal (i.e., a system of representatives of the right cosets of H in G). Then

$$(\phi \uparrow_T G)(g) = [\dot{\phi}(t_i g t_j^{-1}) \mid i, j = 1 \dots k],$$

where $\dot{\phi}(x) = \phi(x)$ for $x \in H$ and the all-zero matrix else, is called the induction of ϕ to G with transversal T . Note that $(\phi \uparrow_T G)(g)$ is a block-permuted matrix, i.e., for all $i = 1 \dots k$ there is exactly one j with $t_i g t_j^{-1} \in H$. If λ has degree one then the induction is monomial.

Finally, recall that a group G is called solvable if there is a sequence $N_i \leq G$, $i = 1 \dots r$, of subgroups, such that

$$\{1\} = N_1 \leq N_2 \leq \dots \leq N_r = G$$

and N_i is normal of prime index in N_{i+1} for $i = 1 \dots r - 1$.

Now we can formulate a coarse version of the recursive decomposition algorithm. The algorithm essentially conjugates a monomial representation to be an induction, which is decomposed along a chain of normal subgroups using a recursion formula for decomposition matrices. For a complete version of this algorithm, including the underlying theorems, we refer the reader to [33].

Algorithm 2 *Given is a monomial representation ϕ of a solvable group G . ϕ shall be decomposed, i.e.,*

$$\phi^A = \rho_1 \oplus \dots \oplus \rho_m,$$

where all ρ_i are irreducible and A is a product of structured sparse matrices.

Case 1: ϕ is not transitive.

1. Decompose ϕ with a permutation P into a direct sum of transitive representations, $\phi^P = \phi_1 \oplus \dots \oplus \phi_\ell$.

2. Recurse with ϕ_i , $i = 1 \dots \ell$, to obtain decomposition matrices B_i .

$A = P \cdot (B_1 \oplus \dots \oplus B_\ell)$ is a decomposition matrix for ϕ .

Case 2: ϕ is transitive.

1. Decompose ϕ with a diagonal matrix D into an induction $\lambda \uparrow_T G$, where λ has degree 1.

2. Recurse with $\lambda \uparrow_T G$ to obtain a decomposition matrix B .

$A = D \cdot B$ is a decomposition matrix for ϕ .

Case 3: $\phi = \lambda \uparrow_T G$, λ is a representation of $H \leq G$, and it exists a normal subgroup $H \leq N \leq G$ of prime index p in G .

1. Decompose ϕ with a monomial matrix M into a double induction $\phi^M = (\lambda \uparrow_{T_1} N) \uparrow_{T_2} G$.

2. Recurse with $\lambda \uparrow_{T_1} N$ to obtain a decomposition matrix B .

$M \cdot A$ is a decomposition matrix for ϕ , where A is given by Theorem 2.

Case 4: $\phi = \lambda \uparrow_T G$, λ is a representation of $H \leq G$, and it exists a normal subgroup $H \not\leq N \leq G$ of prime index p in G .

We omit this step.

Note that at least one of the cases always applies, since G is solvable. We omitted Case 4 since it did not play a role for the examples considered in this paper. The recursion formula for Case 3 is given in Theorem 2. The reader might only want to look at the actual formula for A , omitting the technical details. Obviously, all factors are sparse. For the special case $G = \mathbb{Z}_n = \langle x \mid x^n = 1 \rangle$ and $\phi_1 : x \mapsto [(1, 2, \dots, n), n]$, the permutation P vanishes, $d = 0$, and the formula reduces exactly to the Cooley/Tukey FFT.

Theorem 2 Let $N \leq G$ be a normal subgroup of prime index p with transversal $T = (t^0, t^1, \dots, t^{(p-1)})$. Assume ϕ is a representation of N of degree n with decomposition matrix B such that $\phi^B = \bigoplus_{i=1}^k \rho_i$, where ρ_1, \dots, ρ_j are exactly those among the ρ_i which have an extension $\bar{\rho}_i$ to G . Denote by $d = \deg(\rho_1) + \dots + \deg(\rho_j)$ the entire degree of the extensible ρ_i and set $\bar{\rho} = \bar{\rho}_1 \oplus \dots \oplus \bar{\rho}_j$. Then exists a permutation matrix P such that

$$A = (\mathbf{1}_p \otimes B) \cdot P \cdot \left(\left(\left(\bigoplus_{t \in T} \bar{\rho}(t) \right) (\text{DFT}_p \otimes \mathbf{1}_d) \right) \oplus \mathbf{1}_{p(n-d)} \right)$$

is a decomposition matrix of $\phi \uparrow_T G$.

C. Example 4 in Detail

We work out Steps 1 and Step 2 for Example 4 in greater detail.

Step 1: We choose to compute the mon-mon symmetry of M . Since the matrix M is real, we only consider $\text{MonMon}_2(M)$. The algorithm first constructs the matrix

$$M' = \begin{bmatrix} 1 & -1 & 2 & -2 & 3 & -3 & 4 & -4 \\ -1 & 1 & -2 & 2 & -3 & 3 & -4 & 4 \\ 4 & -4 & 1 & -1 & 2 & -2 & 3 & -3 \\ -4 & 4 & -1 & 1 & -2 & 2 & -3 & 3 \\ 3 & -3 & 4 & -4 & 1 & -1 & 2 & -2 \\ -3 & 3 & -4 & 4 & -1 & 1 & -2 & 2 \\ 2 & -2 & 3 & -3 & 4 & -4 & 1 & -1 \\ -2 & 2 & -3 & 3 & -4 & 4 & -1 & 1 \end{bmatrix}.$$

Any $\{-1, 0, 1\}$ -monomial operation on M (permuting and/or negating rows or columns) is a permutation operation on M' . For example, negating the first column of M can be expressed as the exchange of the first two columns of M' . Now we compute all pairs (L', R') of permutations such that $L'M' = M'R'$ by recursive search. Assume L' maps row one to itself. Then R' must map column one to itself as well, \dots . Assume L' maps row one to row two, then, \dots , etc. As the result we find that $L' = R'$ (which is obvious since the diagonal values of M' are different from the rest so they can only be exchanged with each other) and

$$L' \in \langle (1, 3, 5, 7)(2, 4, 6, 8), (1, 2)(3, 4)(5, 6)(7, 8) \rangle.$$

Translating the permutations on M' back into monomial operations on M , we see that the first permutation is just $[(1, 2, 3, 4), 4]$ and the second is a scalar multiplication with -1 , which we discard since it does not carry any information about M . Hence, we have found the representations $\phi_1 = \phi_2 : x \mapsto [(1, 2, 3, 4), 4]$ of the cyclic group $G = \{1, x, x^2, x^3\}$ as the mon-mon symmetry of M .

Step 2: The remaining task is to decompose the representations $\phi_1 = \phi_2$ using Algorithm 2. Obviously, ϕ_1 is transitive and we apply Case 2 to find $\phi_1 = 1_E \uparrow_T G$, where 1_E is the trivial representation, $1 \mapsto 1$, of the trivial subgroup $E = \{1\} \leq G$, and T is the list $(1, x, x^2, x^3)$ of elements of G . The matrix D is the identity. Next we apply Case 3 using the subgroup $H = \{1, x^2\} \leq G$ and find $(1_E \uparrow_T G)^{[(2,3),4]} = (1_E \uparrow_{T_1} H) \uparrow_{T_2} G$ with $T_1 = (1, x^2)$ and $T_2 = (1, x)$. Now we apply Theorem 2. The matrix $B = \text{DFT}_2$ decomposes $(1_E \uparrow_{T_1} H) : x^2 \mapsto [(1, 2), 2]$ into $\rho_1 : x^2 \mapsto 1$, $\rho_2 : x^2 \mapsto -1$. Both representations can be extended to G through $\bar{\rho}_1 : x \mapsto 1$, and $\bar{\rho}_2 : x \mapsto j$, respectively. In the recursion formula for A , the permutation P vanishes, and evaluating $\bar{\rho} = \bar{\rho}_1 \oplus \bar{\rho}_2$ at the transversal $T_2 = (1, x)$ yields the twiddle factors $\text{diag}(1, 1, 1, j)$. The parameter d equals zero, and we obtain the Cooley/Tukey formula for DFT_4 :

$$A = (\text{DFT}_2 \otimes \mathbf{1}_2) \cdot \text{diag}(1, 1, 1, j) \cdot (\mathbf{1}_2 \otimes \text{DFT}_2) \cdot [(2, 3), 4].$$