Type-IV DCT, DST, and MDCT algorithms with reduced numbers of arithmetic operations

Xuancheng Shao and Steven G. Johnson*

Abstract—**We present algorithms for the type-IV discrete cosine transform (DCT-IV) and discrete sine transform (DST-IV), as well as for the modified discrete cosine transform (MDCT) and its inverse, that achieve a lower count of real multiplications and additions than previously published algorithms, without sacrificing numerical accuracy. Asymptotically, the operation count is reduced from $2N \log_2 N + O(N)$ to $\frac{17}{9} N \log_2 N + O(N)$ for a power-of-two transform size N, and the exact count is strictly lowered for all $N \geq 8$. These results are derived by considering the DCT to be a special case of a DFT of length $8N$, with certain symmetries, and then pruning redundant operations from a recent improved fast Fourier transform algorithm (based on a recursive rescaling of the conjugate-pair split-radix algorithm). The improved algorithms for DST-IV and MDCT follow immediately from the improved count for the DCT-IV.**

Index Terms—**discrete cosine transform; lapped transform; fast Fourier transform; arithmetic complexity**

I. Introduction

In this paper, we present recursive algorithms for type-IV discrete cosine and sine transforms (DCT-IV and DST-IV) and modified discrete cosine transforms (MDCTs), of power-of-two sizes N, that require fewer total real additions and multiplications (herein called *flops*) than previously published algorithms (with an asymptotic reduction of about 6%), without sacrificing numerical accuracy. This work, extending our previous results for small fixed N [1], appears to be the first time in over 20 years that flop counts for the DCT-IV and MDCT have been reduced—although computation times are no longer generally determined by arithmetic counts [2], the question of the minimum number of flops remains of fundamental theoretical interest. Our fast algorithms are based on one of two recently published fast Fourier transform (FFT) algorithms [1], [3], which reduced the operation count for the discrete Fourier transform (DFT) of size N to $\frac{34}{9} N \log_2 N + O(N)$ compared to the (previous best) split-radix algorithm's $4N \log_2 N + O(N)$ [4]–[8]. Given the new FFT, we treat a DCT as an FFT of real-symmetric inputs and eliminate redundant operations to derive the new algorithm; in other work, we applied the same approach to derive improved algorithms for the type-II and type-III DCT and DST [9].

The algorithm for DCT-IV that we present has the same recursive structure as some previous DCT-IV algorithms, but the subtransforms are recursively rescaled in order to eliminate some of the multiplications. This approach reduces the flop count for the DCT-IV from the previous best of $2N \log_2 N + N$

* Department of Mathematics, Massachusetts Institute of Technology, Cambridge MA 01239.

N	previous DCT-IV	New algorithm
8	56	54
16	144	140
32	352	338
64	832	800
128	1920	1838
256	4352	4164
512	9728	9290
1024	21504	20520
2048	47104	44902
4096	102400	97548

TABLE I

Flop counts (real adds + mults) of previous best DCT-IV and our new algorithm

[10]–[16] to:

$$\frac{17}{9} N \log_2 N + \frac{31}{27} N + \frac{2}{9}(-1)^{\log_2 N} \log_2 N - \frac{4}{27}(-1)^{\log_2 N}.$$
(1)

The first savings occur for $N = 8$, and are summarized in Table. I. In order to derive a DCT-IV algorithm from the new FFT algorithm, we simply consider the DCT-IV to be a special case of a DFT with real input of a certain symmetry, and discard the redundant operations. [This should not be confused with algorithms that employ an *unmodified* FFT combined with $O(N)$ pre/post-processing steps to obtain the DCT.] This well-known technique [1], [2], [7]–[9], [17], [18] allows any improvements in the DFT to be immediately translated to the DCT-IV, is simple to derive, avoids cumbersome re-invention of the "same" algorithm for each new trigonometric transform, and (starting with a split-radix FFT) matches the best previous flop counts for every type of DCT and DST. The connection to a DFT of symmetric data can also be viewed as the basic reason why DCT flop counts had not improved for so long: as we review below, the old DCT flop counts can be derived from a split-radix algorithm [15], and the 1968 flop count of split radix was only recently improved upon [1], [3]. There have been many previously published DCT-IV algorithms derived by a variety of techniques, some achieving $2N \log_2 N + N$ flops [10]–[16] and others obtaining larger or unreported flop counts [19]–[22]. Furthermore, exactly the same flop count (1) is now obtained for the type-IV discrete sine transform (DST-IV), since a DST-IV can be obtained from a DCT-IV via flipping the sign of every other input (zero flops, since the sign changes can be absorbed by converting subsequent additions into subtractions or vice versa) and a permutation of the output (zero flops) [12]. Also, in many practical circumstances the output can be scaled by an arbitrary factor (since any scaling

can be absorbed into a subsequent computation); in this case, similar to the well-known savings for a scaled-output size-8 DCT-II in JPEG compression [1], [9], [23], [24], we show that an additional N multiplications can be saved for a scaled-output (or scaled-input) DCT-IV.

Indeed, if we only wished to show that the asymptotic flop count for DCT-IV could be reduced to $\frac{17}{9}N\log_2 N + O(N)$, we could simply apply known algorithms to express a DCT-IV in terms of a real-input DFT (e.g. by reducing it to the DCT-III [10] and thence to a real-input DFT [25]) to immediately apply the $\frac{17}{9}N\log_2 N + O(N)$ flop count for a real-input DFT from our previous paper [1]. However, with FFT and DCT algorithms, there is great interest in obtaining not only the best possible asymptotic constant factor, but also the best possible *exact* count of arithmetic operations. Our result (1) is intended as a new upper bound on this (still unknown) minimum exact count, and therefore we have done our best with the $O(N)$ terms as well as the asymptotic constant.

An important transform closely related to the DCT-IV is an MDCT, which takes $2N$ inputs and produces N outputs, and is designed to be applied to 50%-overlapped blocks of data [26], [27]. Such a "lapped" transform reduces artifacts from block boundaries and is widely used in audio compression [28]. In fact, an MDCT is *exactly* equivalent to a DCT-IV of size N, where the $2N$ inputs have been preprocessed with N additions/subtractions [29]–[32]. This means that the flop count for an MDCT is at most that of a DCT-IV plus N flops. Precisely this technique led to the best previous flop count for an MDCT, $2N\log_2 N + 2N$ [29]–[32]. (There have also been several MDCT algorithms published with larger or unreported flop counts [33]–[42].) It also means that our improved DCT-IV immediately produces an improved MDCT, with a flop count of eq. (1) plus N. Similarly for the inverse MDCT (IMDCT), which takes N inputs to $2N$ outputs and is equivalent to a DCT-IV of size N plus N negations (which should not, we argue, be counted in the flops because they can be absorbed by converting subsequent additions into subtractions).

In the following sections, we first briefly review the new FFT algorithm, previously described in detail [1]. Then, we review how a DCT-IV may be expressed as a special case of a real DFT, and how the new DCT-IV algorithm may be derived by applying the new FFT algorithm and pruning the redundant operations. In doing so, we find it necessary to develop a algorithm for a DCT-III with scaled output. (Previously, we had derived a fast algorithm for a DCT-III based on the new FFT, but only for the case of scaled or unscaled *input* [9].) This DCT-III algorithm follows the same approach of eliminating redundant operations from our new scaled-output FFT (a subtransform of the new FFT) applied to appropriate real-symmetric inputs. We then analyze the flop counts for the DCT-III and DCT-IV algorithms. Finally, we show that this improved DCT-IV immediately leads to improved DST-IV, MDCT, and IMDCT algorithms. We close with some concluding remarks about future directions.

II. REVIEW OF THE NEW FFT

To obtain the new FFT, we used as our starting point a variation called the *conjugate-pair FFT* of the well-known split-radix algorithm. Here, we first review the conjugate-pair FFT, and then briefly summarize how this was modified to reduce the number of flops.

A. Conjugate-pair FFT

The discrete Fourier transform of size N is defined by

$$X_k = \sum_{n=0}^{N-1} x_n \omega_N^{nk}, \qquad (2)$$

where $\omega_N = e^{-\frac{2\pi i}{N}}$ is an Nth primitive root of unity and $k = 0, \ldots, N-1$.

Starting with this equation, the decimation-in-time conjugate-pair FFT [1], [43], a variation on the well-known split-radix algorithm [4]–[7], splits it into three smaller DFTs: one of size $N/2$ of the even-indexed inputs, and two of size $N/4$:

$$X_k = \sum_{n_2=0}^{N/2-1} \omega_{N/2}^{n_2 k} x_{2n_2} + \omega_N^k \sum_{n_4=0}^{N/4-1} \omega_{N/4}^{n_4 k} x_{4n_4+1} + \omega_N^{-k} \sum_{n_4=0}^{N/4-1} \omega_{N/4}^{n_4 k} x_{4n_4-1}. \quad (3)$$

[In contrast, the ordinary split-radix FFT uses x_{4n_4+3} for the third sum (a cyclic shift of x_{4n_4-1}), with a corresponding multiplicative "twiddle" factor of ω_N^{3k}.] This decomposition is repeated recursively until base cases of size $N = 1$ or $N = 2$ are reached. The number of flops required by this algorithm, after certain simplifications (common subexpression elimination and constant folding) and not counting data-independent operations like the computation of ω_N^k, is $4N\log_2 N - 6N + 8$, identical to ordinary split radix [1], [44]–[46].

B. New FFT

Based on the conjugate-pair split-radix FFT from section II-A, a new FFT algorithm with a reduced number of flops can be derived by scaling the subtransforms [1]. We will not reproduce the derivation here, but will simply summarize the results. In particular, the original conjugate-pair split-radix algorithm is split into four mutually recursive algorithms, $\mathbf{newfftS}_N^\ell(x)$ for $\ell = 0, 1, 2, 4$, each of which has the same split-radix structure but computes a DFT scaled by a factor of $1/s_{\ell N,k}$ (defined below), respectively. These algorithms are shown in Algorithm 1, in which the scaling factors are combined with the twiddle factors ω_N^k to reduce the total number of multiplications. In particular, all of the savings occur in $\mathbf{newfftS}_N^1(x)$, while $\mathbf{newfftS}_N^4(x)$ is factorized into a special form to minimize the number of extra multiplications it requires. Here, although $\ell = 0, 1, 2$ are presented in a compact form by a single subroutine $\mathbf{newfftS}_N^\ell(x)$, in practice they would have to be implemented as three separate subroutines in order to exploit the special cases of the multiplicative constants $s_{N,k}/s_{\ell N,k}$, as described in Ref. [1]. For simplicity, we have

Algorithm 1 New FFT algorithm of length N (divisible by 4). The sub-transforms $\mathbf{newfftS}_N^\ell(x)$ for $\ell \neq 0$ are scaled by $s_{\ell N,k}$, respectively, while $\ell = 0$ is the final unscaled DFT ($s_{0,k} = 1$).

---

function $X_{k=0..N-1} \leftarrow \mathbf{newfftS}_N^\ell(x_n)$:
  {computes DFT / $s_{\ell N,k}$, $\ell = 0, 1, 2$}
  $U_{k_2=0...N/2-1} \leftarrow \mathbf{newfftS}_{N/2}^{2\ell}(x_{2n_2})$
  $Z_{k_4=0...N/4-1} \leftarrow \mathbf{newfftS}_{N/4}^1(x_{4n_4+1})$
  $Z'_{k_4=0...N/4-1} \leftarrow \mathbf{newfftS}_{N/4}^1(x_{4n_4-1})$
  for $k = 0$ to $N/4 - 1$ **do**
    $X_k \leftarrow U_k + \left(t_{N,k}Z_k + t_{N,k}^* Z'_k\right) \cdot (s_{N,k}/s_{\ell N,k})$
    $X_{k+N/2} \leftarrow U_k - \left(t_{N,k}Z_k + t_{N,k}^* Z'_k\right) \cdot (s_{N,k}/s_{\ell N,k})$
    $X_{k+N/4} \leftarrow U_{k+N/4}$
      $- i\left(t_{N,k}Z_k - t_{N,k}^* Z'_k\right) \cdot (s_{N,k}/s_{\ell N,k+N/4})$
    $X_{k+3N/4} \leftarrow U_{k+N/4}$
      $+ i\left(t_{N,k}Z_k - t_{N,k}^* Z'_k\right) \cdot (s_{N,k}/s_{\ell N,k+N/4})$
  end for

function $X_{k=0..N-1} \leftarrow \mathbf{newfftS}_N^4(x_n)$:
  {computes DFT / $s_{4N,k}$}
  $U_{k_2=0...N/2-1} \leftarrow \mathbf{newfftS}_{N/2}^2(x_{2n_2})$
  $Z_{k_4=0...N/4-1} \leftarrow \mathbf{newfftS}_{N/4}^1(x_{4n_4+1})$
  $Z'_{k_4=0...N/4-1} \leftarrow \mathbf{newfftS}_{N/4}^1(x_{4n_4-1})$
  for $k = 0$ to $N/4 - 1$ **do**
    $X_k \leftarrow \left[U_k + \left(t_{N,k}Z_k + t_{N,k}^* Z'_k\right)\right] \cdot (s_{N,k}/s_{4N,k})$
    $X_{k+N/2} \leftarrow \left[U_k - \left(t_{N,k}Z_k + t_{N,k}^* Z'_k\right)\right]$
      $\cdot (s_{N,k}/s_{4N,k+N/2})$
    $X_{k+N/4} \leftarrow \left[U_{k+N/4} - i\left(t_{N,k}Z_k - t_{N,k}^* Z'_k\right)\right]$
      $\cdot (s_{N,k}/s_{4N,k+N/4})$
    $X_{k+3N/4} \leftarrow \left[U_{k+N/4} + i\left(t_{N,k}Z_k - t_{N,k}^* Z'_k\right)\right]$
      $\cdot (s_{N,k}/s_{4N,k+3N/4})$
  end for

---

omitted the base cases of the recursion ($N = 1$ and 2) and have not eliminated common subexpressions.

The key aspect of these algorithms is the scale factor $s_{N,k}$, where the subtransforms compute the DFT scaled by $1/s_{\ell N,k}$ for $\ell = 1, 2, 4$. This scale factor is defined for $N = 2^m$ by the following recurrence, where $k_4 = k \mod \frac{N}{4}$:

$$s_{N=2^m,k} = \begin{cases} 1 & \text{for } N \leq 4 \\ s_{N/4,k_4} \cos(2\pi k_4/N) & \text{for } k_4 \leq N/8 \\ s_{N/4,k_4} \sin(2\pi k_4/N) & \text{otherwise} \end{cases} \quad (4)$$

This definition has the properties: $s_{N,0} = 1$, $s_{N,k+N/4} = s_{N,k}$, and $s_{N,N/4-k} = s_{N,k}$. Also, $s_{N,k} > 0$ and decays rather slowly with N: $s_{N,k}$ is $\Omega(N^{\log_4 \cos(\pi/5)})$ asymptotically [1]. When these scale factors are combined with the twiddle factors ω_N^k, we obtain terms of the form

$$t_{N,k} = \omega_N^k \frac{s_{N/4,k}}{s_{N,k}}, \quad (5)$$

which is always a complex number of the form $\pm 1 \pm i \tan \frac{2\pi k}{N}$ or $\pm \cot \frac{2\pi k}{N} \pm i$ and can therefore be multiplied with two fewer

real multiplications than are required to multiply by ω_N^k. We denote the complex conjugate of $t_{N,k}$ by $t_{N,k}^*$. The resulting flop count, for arbitrary complex data x_n, is then reduced from $4N \log_2 N - 6N + 8$ for split radix to $\frac{34}{9} N \log_2 N + O(N)$ [1].

III. FAST DCT-IV FROM NEW FFT

Various forms of discrete cosine transform have been defined, corresponding to different boundary conditions on the transform. The type-IV DCT is defined as a real, linear transformation by the formula:

$$C_k^{\mathrm{IV}} = \sum_{n=0}^{N-1} x_n \cos\left[\frac{\pi}{N}\left(n + \frac{1}{2}\right)\left(k + \frac{1}{2}\right)\right], \quad (6)$$

for N inputs x_n and N outputs C_k^{IV}. The transform can be made orthogonal (unitary) by multiplying with a normalization factor $\sqrt{2/N}$, but for our purposes the unnormalized form is more convenient (and has no effect on the number of operations). We will now derive an algorithm, starting from the new FFT of the previous section, to compute the DCT-IV in terms of a scaled DCT-III and DST-III. These type-III transforms are then treated in following section by a similar method, and lead to our new flop count for the DCT-IV.

In particular, we wish to emphasize in this paper that the DCT-IV (and, indeed, all types of DCT) can be viewed as special cases of the discrete Fourier transform (DFT) with real inputs of a certain symmetry. This viewpoint is fruitful because it means that any FFT algorithm for the DFT leads immediately to a corresponding fast algorithm for the DCT-IV simply by discarding the redundant operations, rather than rederiving a "new" algorithm from scratch. A similar viewpoint has been used to derive fast algorithms for the DCT-II [1], [2], [7], [8], [17], [18], as well as in automatic code-generation for the DCT-IV [1], [2], and has been observed to lead to the minimum known flop count starting from the best known DFT algorithm. Furthermore, because the algorithm is equivalent to an FFT algorithm with certain inputs, it should have the same floating-point error characteristics as that FFT—in this case, the underlying FFT algorithm is simply a rescaling of split radix [1], and therefore inherits the favorable $O(\sqrt{\log N})$ mean error growth and $O(\log N)$ error bounds of the Cooley-Tukey algorithm [47]–[49], unlike at least one other DCT-IV algorithm [12] that has been observed to display $O(\sqrt{N})$ error growth [2].

A. DCT-IV in terms of DFT

Recall that the discrete Fourier transform of size N is defined by eq. (2). In order to derive C_k^{IV} from the DFT formula, one can use the identity $\cos\frac{\pi \ell}{N} =$
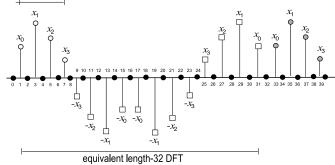
length-4 DCT-IV inputs

equivalent length-32 DFT

Fig. 1. A DCT-IV of length $N = 4$ (open dots x_0, x_1, x_2, x_3) is equivalent to a size $8N = 32$ DFT via interleaving with zeros (black dots) and extending to an odd-even-odd (square dots) periodic (gray dots) sequence.

$\frac{1}{4}\left(\omega_{8N}^{4\ell} - \omega_{8N}^{4N-4\ell} - \omega_{8N}^{4N+4\ell} + \omega_{8N}^{8N-4\ell}\right)$ to write:

$$
\begin{aligned}
C_k^{\mathrm{IV}} &= \sum_{n=0}^{N-1} x_n \cos\left[\frac{\pi}{N}\left(n + \frac{1}{2}\right)\left(k + \frac{1}{2}\right)\right] \\
&= \sum_{n=0}^{N-1} \frac{x_n}{4}\left[\omega_{8N}^{(2n+1)(2k+1)} - \omega_{8N}^{(4N-2n-1)(2k+1)} \right. \\
&\quad \left. - \omega_{8N}^{(4N+2n+1)(2k+1)} + \omega_{8N}^{(8N-2n-1)(2k+1)}\right] \\
&= \sum_{n=0}^{8N-1} \tilde{x}_n \omega_{8N}^{n(2k+1)},
\end{aligned}
\tag{7}
$$

where $\tilde{x}_n$ is a real-even sequence of length $\tilde{N} = 8N$, defined as follows for $0 \le n < N$:

$$
\tilde{x}_{2n+1} = \tilde{x}_{8N-2n-1} = \frac{1}{4}x_n \tag{8}
$$

$$
\tilde{x}_{4N-2n-1} = \tilde{x}_{4N+2n+1} = -\frac{1}{4}x_n \tag{9}
$$

Furthermore, the even-indexed inputs are zeros: $\tilde{x}_{2n} = 0$ for all $0 \le n < 4N$. (The factors of $1/4$ will disappear in the end because they cancel equivalent factors in the subtransforms.)

This is illustrated by an example, for $N = 4$, in Fig. 1. The original four inputs of the DCT-IV are shown as open dots, which are interleaved with zeros (black dots) and extended to an odd-even-odd (square dots) periodic (gray dots) sequence of length $8N = 32$ for the corresponding DFT. Referring to eq. (7), the output of the DCT-IV is given by the first N odd-index outputs of the corresponding DFT (with a scale factor of $1/4$). (The type-IV DCT is distinguished from the other types by the fact that it is even around the left boundary of the original data while it is odd around the right boundary, and the symmetry points fall halfway in between pairs of the original data points.) We will refer, below, to this figure in order to illustrate what happens when an FFT algorithm is applied to this real-symmetirc zero-interleaved data.

B. DCT-IV from DCT/DST-III

For a DCT-IV of size N, our strategy is to directly apply the new FFT algorithm to the equivalent DFT of size $\tilde{N} = 8N$, and to discard the redundant operations from each stage. As it turns out, the sub-transforms after one step of this algorithm are actually scaled-output type-III DCTs and DSTs. This is closely related to a well-known algorithm to express a DCT-IV in terms of a half-size DCT-III and DST-III [10]. In this section, we derive this reduction to a DCT-III for the new FFT algorithm, and then in Sec. IV we derive a new algorithm for the scaled-output DCT-III. The scaled output DST-III algorithm can be easily re-expressed in terms of the DCT-III, as is presented in Sec. IV-C. Here, we define the DCT-III and DST-III, respectively, by the (unnormalized) equations:

$$
C_k^{\mathrm{III}} = \sum_{n=0}^{N-1} x_n \cos\left[\frac{\pi}{N}n\left(k + \frac{1}{2}\right)\right], \tag{10}
$$

$$
S_k^{\mathrm{III}} = \sum_{n=1}^{N} x_n \sin\left[\frac{\pi}{N}n\left(k + \frac{1}{2}\right)\right]. \tag{11}
$$

Starting with the DFT of length $\tilde{N}$ in eq. (2), the new split-radix FFT algorithm splits it into three smaller DFTs: $U_k = \mathrm{dft}(\tilde{x}_{2n_2})$ of size $\tilde{N}/2$, as well as the scaled transforms $Z_k = \frac{1}{s_{\tilde{N}/4,k}}\mathrm{dft}(2\tilde{x}_{4n_4+1})$ and $Z_k' = \frac{1}{s_{\tilde{N}/4,k}}\mathrm{dft}(2\tilde{x}_{4n_4-1})$ of size $\tilde{N}/4$ (including a factor of 2 for convenience). These are combined via:

$$
X_k = U_k + \omega_{\tilde{N}}^k s_{\tilde{N}/4,k}Z_k/2 + \omega_{\tilde{N}}^{-k}s_{\tilde{N}/4,k}Z_k'/2. \tag{12}
$$

Here, where $\tilde{x}_n$ comes from the DCT-IV as in Sec. III-A, the even-indexed elements in $\tilde{x}_n$ are all zero, so $U_k = 0$. Furthermore, by the even symmetry of $\tilde{x}_n$, we have $Z_k' = Z_k^*$ (complex conjugate of Z_k). Thus, we have $C_k^{\mathrm{IV}} = X_{2k+1} = \mathrm{Re}(\omega_{8N}^{2k+1}s_{2N,2k+1}Z_{2k+1})$. We will now show that Z_k is given by combining a DCT-III and a DST-III.

In order to calculate Z_k, we denote for simplicity the inputs of this subtransform by $z_k = 2\tilde{x}_{4k+1}$ for $0 \le k < 2N$. Since Z_k is the output of a real-input DFT of size $2N$, we have $Z_{2N-k} = Z_k^*$. Thus, for any $0 \le k < N/2$, $Z_{2(N-1-k)+1} = Z_{2N-(2k+1)} = Z_{2k+1}^*$. However, there is an additional redundancy in this transform that we must exploit: by inspection of the construction of $\tilde{x}_n$ and by reference to Fig. 2, we see that the inputs z_k are actually a real *anti*-periodic sequence of length N (which becomes periodic when it is extended to length $2N$). We must exploit this symmetry in order to avoid wasting operations.

In particular, by using the anti-periodic symmetry of z_k, we can write the DFT of length $2N$ as a single summation of length N:

$$
\begin{aligned}
Z_{2k+1} &= \frac{1}{s_{2N,2k+1}}\sum_{n=0}^{2N-1}\omega_{2N}^{n(2k+1)}z_n \\
&= \frac{1}{s_{2N,2k+1}}\left(\sum_{n=0}^{N-1}\omega_{2N}^{n(2k+1)}z_n \right. \\
&\quad \left. + \sum_{n=0}^{N-1}\omega_{2N}^{(n+N)(2k+1)}z_{n+N}\right) \\
&= \frac{2}{s_{2N,2k+1}}\sum_{n=0}^{N-1}\omega_{2N}^{n(2k+1)}z_n,
\end{aligned}
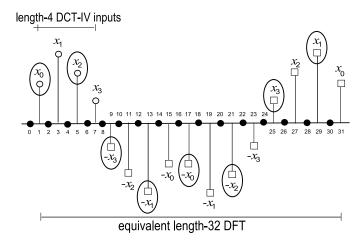$$

length-4 DCT-IV inputs

equivalent length-32 DFT

Fig. 2. The DCT-IV of size 4 (open dots) is computed, in a split-radix conjugate-pair FFT of the $\tilde{N} = 32$ extended data $\tilde{x}_n$ from Fig. 1, via the DFT Z_k of the circled points $\tilde{x}_{4n+1}$, which is an anti-periodic sequence of length $2N = 8$.

Algorithm 2 Fast DCT-IV algorithm in terms of scaled-output DCT-III and DST-III, derived from Algorithm 1 by discarding redundant operations.

---

function $C_{k=0..N-1}^{\mathrm{IV}} \leftarrow \mathbf{newdctIV}_N(x_n)$:
  {computes DCT-IV}
  $w_0 \leftarrow x_0$
  $v_{N/2} \leftarrow x_{N-1}$
  for $k = 1$ to $N/2 - 1$ **do**
    $w_k \leftarrow x_{2k} + x_{2k-1}$
    $v_k \leftarrow x_{2k-1} - x_{2k}$
  end for
  $W_{k=0...N/2-1} \leftarrow \mathbf{newdctIII}_{N/2}^1(w_n)$
  $V_{k=0...N/2-1} \leftarrow \mathbf{newdstIII}_{N/2}^1(v_n)$
  for $k = 0$ to $N/2 - 1$ **do**
    $Z_{2k+1} \leftarrow W_k + iV_k$
    $Z_{2(N-1-k)+1} \leftarrow W_k - iV_k$
  end for
  for $k = 0$ to $N - 1$ **do**
    $C_k^{\mathrm{IV}} \leftarrow \mathrm{Re}\left(\omega_{8N}^{2k+1} s_{2N,2k+1} Z_{2k+1}\right)$
  end for

---

using the facts that $\omega_{2N}^{N(2k+1)} = \omega_2^{(2k+1)} = -1$ and that $z_{n+N} = -z_n$. Then, if we take the real and imaginary parts of the third line above, we obtain precisely a DCT-III and a DST-III, respectively, with outputs scaled by $1/s_{2N,2k+1}$. However, these sub-transforms are actually of size $N/2$, because the symmetry $\omega_{2N}^{(N-n)(2k+1)} = -\omega_{2N}^{-n(2k+1)}$ means that the z_n and z_{N-n} terms merely add or subtract in the input:

$$\mathrm{Re}\, Z_{2k+1} = \frac{2}{s_{2N,2k+1}} \sum_{n=0}^{N-1} \cos\left[\frac{\pi}{N/2} n \left(k + \frac{1}{2}\right)\right] z_n$$

$$= \frac{2z_0 + \sum_{n=1}^{N/2-1} \cos\left[\frac{\pi n\left(k+\frac{1}{2}\right)}{N/2}\right] \cdot 2(z_n - z_{N-n})}{s_{2N,2k+1}} \quad (13)$$

$$\mathrm{Im}\, Z_{2k+1} = -\frac{2}{s_{2N,2k+1}} \sum_{n=0}^{N-1} \sin\left[\frac{\pi}{N/2} n \left(k + \frac{1}{2}\right)\right] z_n$$

$$= \frac{(-1)^k 2z_{N/2} + \sum_{n=1}^{N/2-1} \sin\left[\frac{\pi n\left(k+\frac{1}{2}\right)}{N/2}\right] \cdot 2(z_n + z_{N-n})}{-s_{2N,2k+1}} \quad (14)$$

for any $0 \le k < N/2$. We can define two new sequences w_n $(0 \le n < N/2)$ and v_n $(1 \le n \le N/2)$ of length $N/2$ to be the inputs of this DCT-III and DST-III, respectively:

$$w_0 = 2z_0, \quad w_{n>0} = 2(z_n - z_{N-n}) \quad (15)$$

$$v_{n<N/2} = -2(z_n + z_{N-n}), \quad v_{N/2} = -2z_{N/2}. \quad (16)$$

With this definition of w_n and v_n, we can conclude from eqs. (13–14), and the definition of DCT-III and DST-III, that the real part of Z_{2k+1} is exactly a scaled-output DCT-III of w_n, while the imaginary part of Z_{2k+1} is exactly a scaled-output DST-III of v_n. (The scale factors of ± 2 will disappear in the end: they combine with the 2 in $z_n = 2\tilde{x}_{4n+1}$ to cancel the $1/4$ in the definition of $\tilde{x}_n$.)

Thus, we have shown that the first half of the sequence $Z_{2k+1}(0 \le k < N)$ can be found from a scaled-output DCT-III and DST-III of length $N/2$. The second half of the sequence Z_{2k+1} can be derived by the relation $Z_{2(N-1-k)+1} = Z_{2k+1}^*$ obtained earlier. Given Z_{2k+1}, the output of the original DCT-IV, C_k^{IV}, can be obtained by the formula $C_k^{\mathrm{IV}} = \mathrm{Re}(\omega_{8N}^{2k+1} s_{2N,2k+1} Z_{2k+1})$. This algorithm, in which the computation of z_k has been folded into the computation of w_k and v_k, is presented in Algorithm 2.

In Algorithm 2, $\mathbf{newdctIII}_N^\ell(u)$ calculates the DCT-III of $\{w_n\}$ scaled by a factor of $1/s_{4\ell N,2k+1}$ for $\ell = 0, 1, 2, 4$, and will be presented in Sec. IV. Similarly, $\mathbf{newdstIII}_N^\ell(v)$ calculates the DST-III of $\{v_n\}$ scaled by a factor of $1/s_{4\ell N,2k+1}$ for $\ell = 0, 1, 2, 4$, and will be presented in Sec. IV-C in terms of $\mathbf{newdctIII}_N^\ell(u)$.

If the scale factors $s_{2N,2k+1}$ are removed (set to 1) in Algorithm 2, we recover a decomposition of the DCT-IV in terms of an ordinary (unscaled) DCT-III and DST-III that was first described by Wang [10]. This well-known algorithm yields a flop count exactly the same as previous results: $2N \log_2 N + N$. (Wang obtained a slightly larger count, apparently due to an error in adding his DCT-III and DST-III counts.) The introduction of the scaling factors in Algorithm 2 reduces the flop count by simplifying some of the multiplications in the scaled DCT-III/DST-III compared to their unscaled versions, as will be derived in Sec. IV. Note that, in Algorithm 2, multiplying by $\omega_{8N}^{2k+1} s_{2N,2k+1}$ does not require any more operations than multiplying by ω_{8N}^{2k+1}, because the constant product $\omega_{8N}^{2k+1} s_{2N,2k+1}$ can be precomputed. Let $M_S(N)$ denote the number of flops saved in $\mathbf{newdctIII}_N^1(u)$ compared to the best-known unscaled DCT-III. We shall prove in Sec. IV-C that the same number of operations, $M_S(N)$, can be saved in $\mathbf{newdstIII}_N^1(u)$. Thus, the total number of flops required by Alg. 2 will be $2N \log_2 N + N - 2M_S(N/2)$. The formula for $M_S(N)$ will

be derived in Sec. IV, leading to the final DCT-IV flop count formula given by eq. (1).

IV. DCT-III FROM NEW FFT

The type-III DCT (for a convenient choice of normalization) is defined by eq. (10) for N inputs x_n and N outputs C_k^{III}. We will now follow a process similar to that in the previous section for the DCT-IV: we first express C_k^{III} in terms of a larger DFT of length $4N$, then apply the new FFT algorithm of Sec. II-B, and finally discard the redundant operations to yield an efficient DCT-III algorithm. The resulting algorithm matches the DCT-III flop count of our previous publication [9], which improved upon classic algorithms by about 6% asymptotically. Unlike our previous DCT-III algorithm, however, the algorithm presented here also gives us an efficient *scaled-output* DCT-III, which saves $M_S(N)$ operations over the classic DCT-III algorithms.

A. DCT-III in terms of DFT

In order to derive C_k^{III} from the DFT formula, one can use the identity $\cos\frac{\pi\ell}{N} = \frac{1}{4}\left(\omega_{4N}^{2\ell} - \omega_{4N}^{2N-2\ell} - \omega_{4N}^{2N+2\ell} + \omega_{4N}^{4N-2\ell}\right)$ to write:

$$
\begin{aligned}
C_k^{\mathrm{III}} &= \sum_{n=0}^{N-1} x_n \cos\left[\frac{\pi}{N}n\left(k+\frac{1}{2}\right)\right] \\
&= x_0 + \sum_{n=1}^{N-1}\frac{x_n}{4}\left[\omega_{4N}^{n(2k+1)} - \omega_{4N}^{(2N-n)(2k+1)} \right. \\
&\qquad\left. - \omega_{4N}^{(2N+n)(2k+1)} + \omega_{4N}^{(4N-n)(2k+1)}\right] \\
&= \sum_{n=0}^{4N-1} \tilde{x}_n \omega_{4N}^{n(2k+1)}
\end{aligned}
\tag{17}
$$

where $\tilde{x}_n$ is a real-even sequence of length $\tilde{N} = 4N$, defined as follows for $0 < n < N$:

$$
\tilde{x}_n = \tilde{x}_{4N-n} = \frac{1}{4}x_n \tag{18}
$$

$$
\tilde{x}_{2N-n} = \tilde{x}_{2N+n} = -\frac{1}{4}x_n, \tag{19}
$$

with $\tilde{x}_0 = x_0/2$, $\tilde{x}_N = 0$, $\tilde{x}_{2N} = -x_0/2$, and $\tilde{x}_{3N} = 0$. (Notice that the definitions of $\tilde{N}$ and $\tilde{x}_n$ here are different from those in Sec. III-A.)

This is illustrated by an example, for $N = 4$, in Fig. 3. This figure is very similar to Fig. 1: both of them are even around the points $n = 0$ and $n = \tilde{N}/2$, and are odd around the points $n = \tilde{N}/4$ and $n = 3\tilde{N}/4$. The difference from the DCT-IV is that these points of symmetry/anti-symmetry are now data points of the original sequence, and so the data is not interleaved with zeros as it was for the DCT-IV. We will refer, below, to this figure in order to illustrate what happens when an FFT algorithm is applied to this real-symmetirc data.
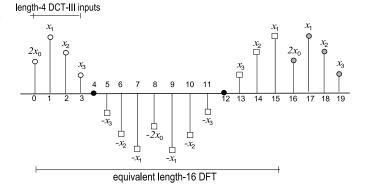.

Fig. 3. A DCT-III of length $N = 4$ (open dots x_0, x_1, x_2, x_3) is equivalent to a DFT of size $4N = 16$ (scaled by a factor of $1/4$), via extending to an odd-even-odd (square dots) periodic (gray dots) sequence and doubling the x_0 term.

B. New (scaled) DCT-III algorithm

In this subsection, we will apply the new FFT algorithm (Alg. 1) to the corresponding DFT for a DCT-III of size N as obtained in Sec. IV-A. This process is similar to what we did in Sec. III-B. We will see that a DCT-III of size N can be calculated by three subtransforms: a DCT-III of size $N/2$, a DCT-III of size $N/4$, and a DST-III of size $N/4$. The resulting algorithm for the DCT-III will have the same recursive structure as in Alg. 1: four mutually recursive subroutines that compute the DCT-III with output scaled by different factors. For use in the DCT-IV algorithm from Sec. III-B, we will actually use only three of these subroutines, because we will only need a scaled-output DCT-III and not the original DCT-III.

When the new FFT algorithm is applied to the sequence $\tilde{x}_n$ of length $\tilde{N} = 4N$ defined by eqs. (18–19), we get three subtransforms: of the sequences $\tilde{x}_{2n_2}$, $\tilde{x}_{4n_4+1}$, and $\tilde{x}_{4n_4-1}$. The DFT of the sequence $\tilde{x}_{2n_2}$ is equivalent to a size-$N/2$ DCT-III of the original even-indexed data x_{2n}, as can be seen from Fig. 3. The subtransforms of $\tilde{x}_{4n_4+1}$ and $\tilde{x}_{4n_4-1}$ have exactly the same properties as the corresponding subtransforms of the DCT-IV as described in Sec. III-B (except that the length of the subtransform $\tilde{x}_{4n_4+1}$ here is N instead of $2N$ as in the DCT-IV case). That is, we denote the DFT of $2\tilde{x}_{4n_4+1}$ by Z_k, and this combines with the DFT $Z_k' = Z_k^*$ of $2\tilde{x}_{4n_4-1}$ to yield a $\mathrm{Re}(\omega_{4N}^{2k+1}Z_{2k+1})$ term in the output as before. And, as before, the inputs of $\tilde{x}_{4n_4+1}$ are anti-periodic with length $N/2$. In consequence, we can apply the result derived in Sec. III-B to conclude that these two subtransforms can be found from a DCT-III of size $N/4$ and a DST-III of size $N/4$. The corresponding inputs of the DCT-III and DST-III, w_n ($0 \le n < N/4$) and v_n ($1 \le n \le N/4$), are defined as follows [compare to eqs. (15–16)]:

$$
w_0 = 2z_0, \; w_{n>0} = 2(z_n - z_{N/2-n}) \tag{20}
$$

$$
v_{n<N/4} = -2(z_n + z_{N/2-n}), \; v_{N/4} = -2z_{N/4}, \tag{21}
$$

where $z_n = 2\tilde{x}_{4n+1}$ for $0 \le n < N/4$. (Again, the factors of 2 will cancel the factor of $1/4$ in the definition of $\tilde{x}_{4n+1}$.)

Therefore, the real part of Z_{2k+1} is the DCT-III of w_n, while the imaginary part is the DST-III of v_n. In summary, a DCT-III of size N can be calculated by a DCT-III of size $N/2$, a DCT-III of size $N/4$, and a DST-III of size $N/4$. Without the scaling factors s, this is equivalent to a decomposition derived by Wang of a DCT-III of size N into a DCT-III and a DCT-IV of size $N/2$ [19], in which the DCT-IV is then decomposed into a DCT-III and a DST-III of size $N/4$ [10].

The above discussion is independent of the scaling factor applied to the output of the transform. So, for the various scale factors in the different subroutines of Alg. 1, we simply scale the DCT-III and DST-III subtransforms in the same way to obtain similar savings in the multiplications (as quantified in the next section). This results in the new DCT-III algorithm presented in Algorithm 3. (The base cases, for $N = 1$ and $N = 2$, are omitted for simplicity.)

Just as in Sec. II-B, although $\ell = 0, 1, 2$ are presented here in a compact form by a single subroutine $\mathbf{newdctIII}_N^\ell (x)$, in practice they would have to be implemented as separate subroutines in order to exploit the special cases of the multiplicative constants $s_{4N,2k+1}/s_{4\ell N,2k+1}$, similar to our FFT [1].

C. DST-III from DCT-III

The DST-III and DCT-III are closely related. In particular, a DST-III can be obtained from a DCT-III, with the same number of flops, by reversing the inputs and multiplying every other output by -1 [9], [12], [50], [51]:

$$S_k^{\mathrm{III}} = \sum_{n=1}^{N} x_n \sin\left[\frac{\pi}{N}\left(k + \frac{1}{2}\right)n\right]$$

$$= (-1)^k \sum_{n=0}^{N-1} x_{N-n} \cos\left[\frac{\pi}{N}\left(k + \frac{1}{2}\right)n\right]. \quad (22)$$

Similarly, one can derive a DST-III in terms of a DCT-III algorithm for any scaling factor. Here, we present the algorithm $\mathbf{newdstIII}_N^\ell (v)$ in terms of $\mathbf{newdctIII}_N^\ell (u)$ for $\ell = 0, 1, 2, 4$. As we can see from Alg. 4, the new DST-III algorithm (with scaled output) has exactly the same operation count as the corresponding DCT-III algorithm (unary negations are not counted in the flops, because they can be absorbed by converting additions into subtractions or vice versa in the preceding DCT computation). This proves our previous assertion that the numbers of operations saved in $\mathbf{newdctIII}_N^1 (u)$ and $\mathbf{newdstIII}_N^1 (v)$, compared to the known unscaled algorithms, are both the same number $M_S(N)$.

V. FLOP COUNTS FOR DCT-III/IV

First, we will show that Alg. 3 gives the best previous flop count $T(N) = 2N \log_2 N - N + 1$ for the DCT-III if the scaling factor s is set to 1. Inspection of Algorithm 2 yields a flop count $4N - 2 + 2T(N/2)$ for the DCT-IV, and substituting $T(N)$ gives the previous best flop count of $2N \log_2 N + N$ for the DCT-IV. Then, we will analyze how many operations are *saved* when the scaling factors are included.

If $s = 1$, we can see from Alg. 3 that a DCT-III of size N is decomposed into a DCT-III of size $N/2$, a DCT-III of size

---

Algorithm 3 New DCT-III algorithm of length N. The sub-transforms $\mathbf{newdctIII}_N^\ell (x)$ for $\ell \neq 0$ are scaled by $s_{4\ell N,2k+1}$, respectively, while $\ell = 0$ is the final unscaled DFT ($s_{0,2k+1} = 1$).

function $C_{k=0..N-1}^{\mathrm{III}} \leftarrow \mathbf{newdctIII}_N^\ell (x_n)$:
 {computes DCT-III / $s_{4\ell N,2k+1}$ for $\ell = 0, 1, 2$}
 $w_0 \leftarrow x_1$
 $v_{N/4} \leftarrow x_{N-1}$
 for $k = 1$ to $N/4 - 1$ **do**
  $w_k \leftarrow x_{4k+1} + x_{4k-1}$
  $v_k \leftarrow x_{4k-1} - x_{4k+1}$
 end for
 $U_{k_2=0...N/2-1} \leftarrow \mathbf{newdctIII}_{N/2}^{2\ell} (x_{2n_2})$
 $W_{k_4=0...N/4-1} \leftarrow \mathbf{newdctIII}_{N/4}^1 (w_{n_4})$
 $V_{k_4=0...N/4-1} \leftarrow \mathbf{newdstIII}_{N/4}^1 (v_{n_4})$
 for $k = 0$ to $N/4 - 1$ **do**
  $Z_{2k+1} \leftarrow W_k + iV_k$
  $Z_{N-2k-1} \leftarrow W_k - iV_k$
 end for
 for $k = 0$ to $N/2 - 1$ **do**
  $C_k^{\mathrm{III}} \leftarrow U_k + \mathrm{Re}\left(t_{4N,2k+1} Z_{2k+1}\right) \frac{s_{4N,2k+1}}{s_{4\ell N,2k+1}}$
  $C_{N-k-1}^{\mathrm{III}} \leftarrow U_k - \mathrm{Re}\left(t_{4N,2k+1} Z_{2k+1}\right) \frac{s_{4N,2k+1}}{s_{4\ell N,2k+1}}$
 end for

function $C_{k=0..N-1}^{\mathrm{III}} \leftarrow \mathbf{newdctIII}_N^4 (x_n)$:
 {computes DCT-III / $s_{16N,2k+1}$}
 $w_0 \leftarrow x_1$
 $v_{N/4} \leftarrow x_{N-1}$
 for $k = 1$ to $N/4 - 1$ **do**
  $w_k \leftarrow x_{4k+1} + x_{4k-1}$
  $v_k \leftarrow x_{4k-1} - x_{4k+1}$
 end for
 $U_{k_2=0...N/2-1} \leftarrow \mathbf{newdctIII}_{N/2}^2 (x_{2n_2})$
 $W_{k_4=0...N/4-1} \leftarrow \mathbf{newdctIII}_{N/4}^1 (w_{n_4})$
 $V_{k_4=0...N/4-1} \leftarrow \mathbf{newdstIII}_{N/4}^1 (v_{n_4})$
 for $k = 0$ to $N/4 - 1$ **do**
  $Z_{2k+1} \leftarrow W_k + iV_k$
  $Z_{N-2k-1} \leftarrow W_k - iV_k$
 end for
 for $k = 0$ to $N/2 - 1$ **do**
  $C_k^{\mathrm{III}} \leftarrow [U_k + \mathrm{Re}\left(t_{4N,2k+1} Z_{2k+1}\right)] \frac{s_{4N,2k+1}}{s_{16N,2k+1}}$
  $C_{N-k-1}^{\mathrm{III}} \leftarrow [U_k - \mathrm{Re}\left(t_{4N,2k+1} Z_{2k+1}\right)] \frac{s_{4N,2k+1}}{s_{16N,2N+2k+1}}$
 end for

Algorithm 4 scaled-output DST-III algorithm of size N, based on the scaled-output DCT-III algorithm which is presented in Sec. IV, with the same operation count.

function $S_{k=0..N-1}^{\mathrm{III}} \leftarrow \mathbf{newdstIII}_N^\ell (x_n)$:
 {computes DST-III /$s_{4\ell N,2k+1}$}
 for $k = 0$ to $N - 1$ **do**
  $w_k \leftarrow x_{N-k}$
 end for
 $C_{k=0...N-1}^{\mathrm{III}} \leftarrow \mathbf{newdctIII}_N^\ell (w_n)$
 for $k = 0$ to $N - 1$ **do**
  $S_k^{\mathrm{III}} \leftarrow (-1)^k C_k^{\mathrm{III}}$
 end for

$N/4$ and a DST-III of size $N/4$, and all four of our recursive subroutines are identical (they only differed by s factors). In addition, $2(N/4-1)$ flops are required to obtain the sequences w_k and v_k, and $5N/2$ flops are required to obtain the output of the DCT-III from the outputs of the subtransforms. Therefore, we obtain the recurrence relation for $T(N)$:

$$T(N) = T(N/2) + 2T(N/4) + 3N - 2. \qquad (23)$$

The initial conditions for $T(N)$ can be determined easily. If $N = 1$, $y_0 = x_0$. Therefore, $T(1) = 0$. If $N = 2$, $y_0 = x_0 + x_1/\sqrt{2}$ and $y_1 = x_0 - x_1/\sqrt{2}$. Therefore, $T(2) = 3$. Solving eq. (23) with these initial conditions, we immediately obtain the following result:

$$T(N) = 2N \log_2 N - N + 1. \qquad (24)$$

This flop count is the same as the previous best flop count for DCT-III algorithms [7], [10]–[13], [17], [21], [51]–[56] prior to our work [1], [9].

Since our DCT-III algorithm without scaling factors (i.e. with $s = 1$) obtains the same number of flops as the best previous DCT-III algorithms, it only remains to determine how many operations are *saved* by including the scale factors. We now analyze this count of saved flops by solving the appropriate recurrence relations. In particular, let $M(N)$, $M_S(N)$, $M_{S2}(N)$ and $M_{S4}(N)$ (where N is a power of 2) denote the number of operation saved (or spent, if negative) in $\mathbf{newdctIII}_N^\ell(x)$ for $\ell = 0, 1, 2, 4$, respectively, compared to the corresponding unscaled DCT-III algorithm.

First, let us derive the recurrence relations for $M(N)$ and so on, similar to the analysis of Alg. 1 [1]. The number of flops saved in $\mathbf{newdctIII}_N^\ell(x)$ is the sum of the flops saved in the subtransforms and the number of flops saved in the loop to calculate the final results C_k^{III}. In $\mathbf{newdctIII}_N^0(x)$, $5 \cdot \frac{N}{2}$ flops are required in the loop, as in the old unscaled algorithm. In $\mathbf{newdctIII}_N^1(x)$, only $4 \cdot \frac{N}{2}$ flops are needed since either the real part or the imaginary part of $t_{4N,2k+1}$ is 1. Thus, $N/2$ flops in the loop are saved for $\ell = 1$. In $\mathbf{newdctIII}_N^2(x)$, $5 \cdot \frac{N}{2}$ flops are again required in the loop. (In contrast, for Alg. 1 the $\ell = 2$ case required two more multiplications than the $\ell = 0$ case because of the $k = 0$ term [1], which is not present here because $2k + 1 \neq 0$.) In $\mathbf{newdctIII}_N^4(x)$, $6 \cdot \frac{N}{2}$ flops are required in the loop since $s_{16N,2k+1} \neq s_{16N,2k+1+2N}$ and hence we must multiply the two scale factors separately. This means that we *spend* $N/2$ extra multiplications in the $\ell = 4$ case, which is counted as a negative term in M_{S4}. Thus, we have the following relations:

$$M(N) = M(N/2) + 2M_S(N/4)$$
$$M_S(N) = M_{S2}(N/2) + 2M_S(N/4) + N/2$$
$$M_{S2}(N) = M_{S4}(N/2) + 2M_S(N/4)$$
$$M_{S4}(N) = M_{S2}(N/2) + 2M_S(N/4) - N/2. \qquad (25)$$

We next determine the number of flops saved (if any) for the base cases, $N = 1$ and $N = 2$. When $N = 1$, the unscaled algorithm computes the output $y_0 = x_0$, while the algorithms $\mathbf{newdctIII}_N^\ell(x)$ calculate $y_0 = (1/s_{4\ell,1})x_0$

which requires the same number of flops for $\ell < 2$ and one more multiplication for $\ell \geq 2$:

$$M(1) = M_S(1) = 0,$$
$$M_{S1}(1) = M_{S4}(1) = -1. \qquad (26)$$

When $N = 2$, we obtain scale factors $s_{4\ell N,2k+1} = s_{8\ell,2k+1}$. For $\ell < 4$, $s_{8\ell,1} = s_{8\ell,3}$, while $s_{8\ell,1} \neq s_{8\ell,3}$ for $\ell = 4$. The unscaled algorithm calculates the output $y_{0,1} = x_0 \pm \sqrt{1/2}x_1$, where 3 flops are required. For $\ell = 0, 1, 2$, the algorithms $\mathbf{newdctIII}_N^\ell(x)$ calculates $y_{0,1} = (x_0 \pm \sqrt{1/2}x_1)/s_{8\ell,1}$, which requires 3 flops for $\ell = 0$ (where $s = 1$), 3 flops for $\ell = 1$ (where $s = 1/\sqrt{2}$ and cancels one of the constants), and 4 flops for $\ell = 2$. For $\ell = 4$, $\mathbf{newdctIII}_N^4(x)$ calculates $y_0 = (x_0 + \frac{1}{\sqrt{2}}x_1)/s_{32,1}$ and $y_1 = (x_0 - \frac{1}{\sqrt{2}}x_1)/s_{32,3}$, where 5 flops are required. Thus, we have

$$M(2) = M_S(2) = 0,$$
$$M_{S2}(2) = -1,$$
$$M_{S4}(2) = -2. \qquad (27)$$

With these base cases, one can solve the recurrences (25) by standard generating-function methods [57] to obtain:

$$M_S(N) = \frac{1}{9}N \log_2 N - \frac{1}{27}N + \frac{1}{9}(-1)^{\log_2 N} \log_2 N + \frac{1}{27}(-1)^{\log_2 N}. \qquad (28)$$

Recall from Sec. III-B that $2M_S(N/2)$ flops can be saved in the new DCT-IV algorithm compared to the best previous algorithms, resulting in a total flop count of $2N \log_2 N + N - 2M_S(N/2)$. This gives the DCT-IV flop count in eq. (1) for Algorithm 2. This expression for the flop count of the new DCT-IV algorithm matches the results that were derived by automatic code generation for small N [1], as expected.

In general, as was discussed in our other work on the DCT-II/III [9], the number of multiplications may change depending upon the normalization chosen. For the DCT-IV, a common normalization choice is to multiply by $\sqrt{2/N}$, which makes the transform unitary, but this does not change the number of flops because the normalization can be absorbed into the $\omega_{8N}^{2k+1} s_{2N,2k+1}$ factor (which is $\neq 1$ for all k). On the other hand, if one is able to scale every output of the DCT-IV individually, for example if the scale factor can be absorbed into a subsequent computational step, then the best choice in the present algorithm seems to be to scale by $1/s_{8N,2k+1}$. This choice of scale factor will transform $\omega_{8N}^{2k+1} s_{2N,2k+1}$ into $t_{8N,2k+1}$ in Algorithm 2, which can be multiplied in one fewer multiplication, saving N multiplications overall. Similarly, one can save N multiplications for a scaled-*input*, unscaled-output DCT-IV, since the scaled-output DCT-IV can be transformed into a scaled-input DCT-IV by network transposition [58] without changing the number of flops [9].

VI. DST-IV FROM DCT-IV

The (unnormalized) DST-IV is defined as:

$$S_k^{\mathrm{IV}} = \sum_{n=0}^{N-1} x_n \sin\left[\frac{\pi}{N}\left(n + \frac{1}{2}\right)\left(k + \frac{1}{2}\right)\right] \qquad (29)$$

for $k = 0, \ldots, N-1$. Although we could derive fast algorithms for S_k^{IV} directly by treating it as a DFT of length $8N$ with odd

symmetry, interleaved with zeros, and discarding redundant operations similar to above, it turns out there is a simpler technique. The DST-IV is *exactly* equivalent to a DCT-IV in which the outputs are reversed and every other input is multiplied by -1 (or vice versa) [12]:

$$S_{N-1-k} = 2 \sum_{n=0}^{N-1} (-1)^n x_n \cos\left[\frac{\pi}{N}\left(n+\frac{1}{2}\right)\left(k+\frac{1}{2}\right)\right] \tag{30}$$

for $k = 0, \ldots, N-1$. It therefore follows that a DST-IV can be computed with the same number of flops as a DCT-IV of the same size, assuming that multiplications by -1 are free—the reason for this is that sign flips can be absorbed at no cost by converting additions into subtractions or vice versa in the subsequent algorithmic steps. Therefore, our new flop count (1) immediately applies to the DST-IV.

VII. MDCT FROM DCT-IV

In this section, we will present a new modified DCT (MDCT) algorithm in terms of our new DCT-IV algorithm with an improved flop count compared to the best previously published counts. The key fact is that the best previous flop count for an MDCT of $2N = 2^m$ inputs and N outputs was obtained by reducing the problem to a DCT-IV plus N extra additions [29]–[32]. Therefore, our improved DCT-IV algorithm immediately gives an improved MDCT. Similarly for the inverse MDCT, except that in that case no extra additions are required.

An MDCT of length "N" has $2N$ inputs x_n ($0 \le n < 2N$) and N outputs C_k^{M} ($0 \le k < N$) defined by the formula (not including normalization factors):

$$C_k^{\mathrm{M}} = \sum_{n=0}^{2N-1} x_n \cos\left[\frac{\pi}{N}\left(n+\frac{1}{2}+\frac{N}{2}\right)\left(k+\frac{1}{2}\right)\right]. \tag{31}$$

This is "inverted" by the inverse MDCT (IMDCT), which takes N inputs C_k^{M} and gives $2N$ outputs y_n, defined by (again not including normalization):

$$y_n = \sum_{k=0}^{N-1} C_k^{\mathrm{M}} \cos\left[\frac{\pi}{N}\left(n+\frac{1}{2}+\frac{N}{2}\right)\left(k+\frac{1}{2}\right)\right]. \tag{32}$$

These transforms are designed to operate on consecutive 50%-overlapping blocks of data, and when the IMDCTs of subsequent blocks are added in their overlapping halves the resulting "time-domain aliasing cancellation" (TDAC) yields the original data x_n [26], [27]. The MDCT is widely used in audio compression, where the overlapping reduces artifacts from the block boundaries [28].

The MDCT and IMDCT can be trivially re-expressed in terms of a DCT-IV of size N [29]–[32]. Let us define

$$\Xi_n = \cos\left[\frac{\pi}{N}\left(n+\frac{1}{2}\right)\left(k+\frac{1}{2}\right)\right],$$

which has the symmetry $\Xi_{2N+n} = \Xi_{2N-1-n} = -\Xi_n$. In terms of Ξ_n, the MDCT becomes

$$\begin{aligned}
C_k^{\mathrm{M}} &= \sum_{n=0}^{2N-1} \Xi_{\frac{N}{2}+n} x_n \\
&= \sum_{n=0}^{\frac{N}{2}-1} \left(\Xi_{\frac{N}{2}+n} x_n + \Xi_{\frac{3N}{2}-1-n} x_{N-1-n} \right. \\
&\qquad \left. + \Xi_{2N-1-n} x_{\frac{3N}{2}-1-n} + \Xi_{2N+n} x_{\frac{3N}{2}+n} \right) \\
&= \sum_{n=0}^{\frac{N}{2}-1} \Xi_{\frac{N}{2}+n} \left(x_n - x_{N-1-n} \right) \\
&\quad - \sum_{n=0}^{\frac{N}{2}-1} \Xi_n \left(x_{\frac{3N}{2}-1-n} + x_{\frac{3N}{2}+n} \right) \\
&= \sum_{n=N/2}^{N-1} \Xi_n \left(x_{n-N/2} - x_{\frac{3N}{2}-1-n} \right) \\
&\quad - \sum_{n=0}^{N/2-1} \Xi_n \left(x_{\frac{3N}{2}-1-n} + x_{\frac{3N}{2}+n} \right).
\end{aligned}$$

But the final summation is simply a DCT-IV of the sequence $\tilde{x}_n$ defined by $\tilde{x}_n = -(x_{\frac{3N}{2}-1-n} + x_{\frac{3N}{2}+n})$ for $0 \le n < \frac{N}{2}$ and $\tilde{x}_n = x_{n-N/2} - x_{\frac{3N}{2}-1-n}$ for $\frac{N}{2} \le n < N$. Therefore, given any algorithm for a DCT-IV, the MDCT can be computed with at most N extra additions. (Here, we are not counting multiplication by -1, because negations can be absorbed by converting additions into subtractions and vice versa in subsequent computational steps.) Since the previous best flop count for the DCT-IV was $2N \log_2 N + N$ flops, this led to a flop count of $2N \log_2 N + 2N$ for the MDCT [29]–[32]. Instead, we can use our new algorithm for the DCT-IV to immediately reduce this flop count for the MDCT to eq. (1)+N:

$$\frac{17}{9} N \log_2 N + \frac{58}{27} N + \frac{2}{9}(-1)^{\log_2 N} \log_2 N - \frac{4}{27}(-1)^{\log_2 N}. \tag{33}$$

The IMDCT requires almost no manipulation: it is already in the form of a DCT-IV, except that we are evaluating the DCT-IV beyond the "end" of the inputs. Since a DCT-IV corresponds to anti-symmetric data as in Fig. 1, this just means that we compute the DCT-IV and obtain the IMDCT by storing the outputs and their mirror image (multiplied by -1), shifted by $N/2$. So, the flop count for the IMDCT is exactly the same as the flop count for the DCT-IV, not counting negations. Any overall negation of an output can be eliminated by converting a preceding addition to a subtraction (or changing the sign of a preceding constant), but some of the (redundant) IMDCT outputs are needed with both signs, which seems to imply that an explicit negation is required. The latter negations are easily eliminated in practice, however: an IMDCT is always followed in practice by adding overlapping IMDCT blocks to achieve TDAC, so the negations simply mean that some of these additions are converted into subtractions.

VIII. CONCLUDING REMARKS

We have derived new algorithms for the DCT-IV, DST-IV, MDCT, and IMDCT that reduce the flops for a size

$N = 2^m$ from $2N \log_2 N + O(N)$ to $\frac{17}{9} N \log_2 N + O(N)$, representing the first improvement in their flop counts for many years and stemming from similar developments for the DFT [1], [3]. We do not claim that these flop counts are the best possible, although we are not currently aware of any way to obtain further reductions in either the leading coefficient or in our $O(N)$ terms. It is possible that further gains could be made by extending our recursive rescaling technique to greater generality, for example. However, we believe that such investigations will be most easily carried out in the context of the DFT, since FFT algorithms (in terms of complex exponentials) are typically much easier to work with than fast DCT algorithms (in terms of real trigonometry), and any improved FFT algorithm will immediately lead to similar gains for DCTs (and vice versa: any improved DCT leads to an improved FFT) [8]. Another open question is whether these new algorithms will lead to practical gains in performance on real computers. This is a complicated and somewhat ill-defined question, because performance characteristics vary between machines and depend strongly on many factors besides flop counts—any simple algorithms like the ones presented here require extensive restructuring to make them efficient on real computers, just as classic split-radix does not perform well without modification [2]. On the other hand, for small fixed N where straight-line (unrolled) hard-coded kernels are often employed in audio and image processing (where the block size is commonly fixed), we have demonstrated that automatic code-generation techniques (given only the new FFT) can produce efficient DCT-IV (and MDCT, etc.) kernels attaining the new operation counts, and that the performance is sometimes improved at least slightly [1].

Acknowledgements

References

[1] S. G. Johnson, M. Frigo, A modified split-radix FFT with fewer arithmetic operations, IEEE Trans. Signal Processing 55 (1) (2007) 111–119.

[2] M. Frigo, S. G. Johnson, The design and implementation of FFTW3, Proc. IEEE 93 (2) (2005) 216–231.

[3] T. Lundy, J. Van Buskirk, A new matrix approach to real FFTs and convolutions of length 2^k, Computing 80 (1) (2007) 23–45.

[4] R. Yavne, An economical method for calculating the discrete Fourier transform, in: Proc. AFIPS Fall Joint Computer Conf., Vol. 33, 1968, pp. 115–125.

[5] P. Duhamel, H. Hollmann, Split-radix FFT algorithm, Electron. Lett. 20 (1) (1984) 14–16.

[6] J. B. Martens, Recursive cyclotomic factorization—a new algorithm for calculating the discrete Fourier transform, IEEE Trans. Acoust., Speech, Signal Processing 32 (4) (1984) 750–761.

[7] M. Vetterli, H. J. Nussbaumer, Simple FFT and DCT algorithms with reduced number of operations, Signal Processing 6 (4) (1984) 267–278.

[8] P. Duhamel, M. Vetterli, Fast Fourier transforms: a tutorial review and a state of the art, Signal Processing 19 (1990) 259–299.

[9] X. Shao, S. G. Johnson, Type-II/III DCT/DST algorithms with reduced number of arithmetic operations, arXiv.org e-Print archive (2007) arXiv:cs.DS/0703150.
URL http://arxiv.org/abs/cs.DS/0703150

[10] Z. Wang, On computing the discrete Fourier and cosine transforms, IEEE Trans. Acoust., Speech, Signal Processing 33 (4) (1985) 1341–1344.

[11] N. Suehiro, M. Hatori, Fast algorithms for the DFT and other sinusoidal transforms, IEEE Trans. Acoust., Speech, Signal Processing 34 (3) (1986) 642–644.

[12] S. C. Chan, K. L. Ho, Direct methods for computing discrete sinusoidal transforms, IEE Proceedings F 137 (6) (1990) 433–442.

[13] C. W. Kok, Fast algorithm for computing discrete cosine transform, IEEE Trans. Signal Processing 45 (3) (1997) 757–760.

[14] V. Britanak, K. R. Rao, The fast generalized discrete Fourier transforms: a unified approach to the discrete sinusoidal transforms computation, Sig. Proc. 79 (1999) 135–150.

[15] G. Plonka, M. Tasche, Split-radix algorithms for discrete trigonometric transforms, online at http://www.uni-duisburg.de/FB11/STAFF/PLONKA/ps.html (2002).

[16] V. Britanak, The fast DCT-IV/DST-IV computation via the MDCT, Sig. Proc. 83 (2003) 1803–1813.

[17] P. Duhamel, Implementation of "split-radix" FFT algorithms for complex, real, and real-symmetric data, IEEE Trans. Acoust., Speech, Signal Processing 34 (2) (1986) 285–295.

[18] R. Vuduc, J. Demmel, Code generators for automatic tuning of numerical kernels: experiences with FFTW, in: Proc. Semantics, Application, and Implementation of Code Generators Workshop, Montreal, 2000.

[19] Z. Wang, Fast algorithms for the discrete W transform and for the discrete Fourier transform, IEEE Trans. Acoust., Speech, Signal Processing 32 (4) (1984) 803–816.

[20] N. R. Murthy, M. N. S. Swamy, On the algorithms for the computation of even discrete cosine transform-2 (EDCT-2) of real sequences, IEEE Trans. Circ. Syst. 37 (5) (1990) 625–627.

[21] M. Püschel, J. M. F. Moura, The algebraic approach to the discrete cosine and sine transforms and their fast algorithms, SIAM J. Computing 32 (5) (2003) 1280–1316.

[22] G. Plonka, M. Tasche, Fast and numerically stable algorithms for discrete cosine transforms, Lin. Algebra and its Appl. 394 (2005) 309–345.

[23] Y. Arai, T. Agui, M. Nakajima, A fast DCT-SQ scheme for images, Trans. IEICE 71 (11) (1988) 1095–1097.

[24] W. B. Pennebaker, J. L. Mitchell, JPEG Still Image Data Compression Standard, Van Nostrand Reinhold, New York, 1993.

[25] M. J. Narasimha, A. M. Peterson, On the computation of the discrete cosine transform, IEEE Trans. Commun. 26 (6) (1978) 934–936.

[26] J. P. Princen, A. B. Bradley, Analysis/synthesis filter bank design based on time domain aliasing cancellation, IEEE Trans. Acoust., Speech, Signal Processing 34 (5) (1986) 1153–1161.

[27] A. W. Johnson, A. B. Bradley, Adaptive transform coding incorporating time domain aliasing cancellation, Speech Commun. 6 (1987) 299–308.

[28] T. Painter, A. Spanias, Perceptual coding of digital audio, Proc. IEEE 88 (4) (2000) 451–515.

[29] H. S. Malvar, Lapped transform for efficient transform/subband coding, IEEE Trans. Acoust., Speech, Signal Processing 38 (1990) 969–978.

[30] H. S. Malvar, Fast algorithm for modulated lapped transform, Electronic Letters 27 (9) (1991) 775–776.

[31] D. Y. Chan, J. F. Yang, S. Y. Chen, Regular implementation algorithms of time domain aliasing cancellation, IEE Proc. Vision Image Signal Processing 143 (6) (1996) 387–392.

[32] C.-M. Liu, W.-C. Lee, A unified fast algorithm for cosine modulated filter banks in current audio coding standards, J. Audio Engineering Soc. 47 (12) (1999) 1061–1075.

[33] H. S. Malvar, Extended lapped transforms: properties, applications, and fast algorithms, IEEE Trans. Sig. Proc. 40 (11) (1992) 2703–2714.

[34] T. Sporer, K. Brandenburg, B. Edler, The use of multirate filter banks for coding of high quality digital audio, in: Proc. 6th European Sig. Processing Conf. (EUSIPCO), Vol. 1, Amsterdam, 1992, pp. 211–214.

[35] H.-C. Chiang, J.-C. Liu, Regressive implementations for the forward and inverse MDCT in MPEG audio coding, IEEE Signal Processing Letters 3 (4) (1996) 116–118.

[36] Y.-H. Fan, V. K. Madisetti, R. M. Mersereau, On fast algorithms for computing the inverse modified discrete cosine transform, IEEE Signal Processing Letters 6 (3) (1999) 61–64.

[37] S.-W. Lee, Improved algorithm for efficient computation of the forward and backward MDCT in MPEG audio coder, IEEE Trans. Circuits and Systems II: Analog and Digital Signal Processing 48 (10) (2001) 990–994.

[38] C. Jing, H.-M. Tai, Fast algorithm for computing modulated lapped transform, Electronic Letters 37 (12) (2001) 796–797.

[39] V. Britanak, K. R. Rao, A new fast algorithm for the unified forward and inverse MDCT/MDST computation, Sig. Proc. 82 (2002) 433–459.

[40] M.-H. Cheng, Y.-H. Hsu, Fast IMDCT and MDCT algorithms – a matrix approach, IEEE Trans. Sig. Proc. 51 (1) (2003) 221–229.

[41] V. Nikolajevic, G. Fettweis, Computation of forward and inverse MDCT using Clenshaw's recurrence formula, IEEE Trans. Sig. Proc. 51 (5) (2003) 1439–1444.

[42] C.-H. Chen, B.-D. Liu, J.-F. Yang, Recursive architectures for realizing modified discrete cosine transform and its inverse, IEEE Trans. Circuits Syst. II 50 (1) (2003) 38–45.

[43] I. Kamar, Y. Elcherif, Conjugate pair fast Fourier transform, Electron. Lett. 25 (5) (1989) 324–325.

[44] R. A. Gopinath, Comment: Conjugate pair fast Fourier transform, Electron. Lett. 25 (16) (1989) 1084.

[45] H.-S. Qian, Z.-J. Zhao, Comment: Conjugate pair fast Fourier transform, Electron. Lett. 26 (8) (1990) 541–542.

[46] A. M. Krot, H. B. Minervina, Comment: Conjugate pair fast Fourier transform, Electron. Lett. 28 (12) (1992) 1143–1144.

[47] W. M. Gentleman, G. Sande, Fast Fourier transforms—for fun and profit, Proc. AFIPS 29 (1966) 563–578.

[48] J. C. Schatzman, Accuracy of the discrete Fourier transform and the fast Fourier transform, SIAM J. Scientific Computing 17 (5) (1996) 1150–1166.

[49] M. Tasche, H. Zeuner, Handbook of Analytic-Computational Methods in Applied Mathematics, CRC Press, Boca Raton, FL, 2000, Ch. 8, pp. 357–406.

[50] Z. Wang, A fast algorithm for the discrete sine transform implemented by the fast cosine transform, IEEE Trans. Acoust., Speech, Signal Processing 30 (5) (1982) 814–815.

[51] P. Lee, F.-Y. Huang, Restructured recursive DCT and DST algorithms, IEEE Trans. Signal Processing 42 (7) (1994) 1600–1609.

[52] B. G. Lee, A new algorithm to compute the discrete cosine transform, IEEE Trans. Acoust., Speech, Signal Processing 32 (6) (1984) 1243–1245.

[53] H. S. Hou, A fast algorithm for computing the discrete cosine transform, IEEE Trans. Acoust., Speech, Signal Processing 35 (10) (1987) 1455–1461.

[54] F. Arguello, E. L. Zapata, Fast cosine transform based on the successive doubling method, Electronic letters 26 (19) (1990) 1616–1618.

[55] J. Takala, D. Akopian, J. Astola, J. Saarinen, Constant geometry algorithm for discrete cosine transform, IEEE Trans. Signal Processing 48 (6) (2000) 1840–1843.

[56] M. Püschel, Cooley–Tukey FFT like algorithms for the DCT, in: Proc. IEEE Int'l Conf. Acoustics, Speech, and Signal Processing, Vol. 2, Hong Kong, 2003, pp. 501–504.

[57] D. E. Knuth, Fundamental Algorithms, 3rd Edition, Vol. 1 of The Art of Computer Programming, Addison-Wesley, 1997.

[58] R. E. Crochiere, A. V. Oppenheim, Analysis of linear digital networks, Proc. IEEE 63 (4) (1975) 581–595.