Exploiting Symmetries of Small Prime-Sized DFTs

Doru Thom Popovici[1], Devangi Parikh[2], Daniele G. Spampinato[3], Tze Meng Low[3]
[1]Lawrence Berkeley National Lab, CA, USA
[2]University of Texas at Austin, TX, USA
[3]Carnegie Mellon University, PA, USA
dtpopovici@lbl.gov, dnp@cs.utexas.edu
{spampinato, lowt}@cmu.edu

Various applications from fields like quantum mechanics, material sciences, and machine learning require computing the discrete Fourier transforms on prime-sized inputs. Typically, this computation is done as a cyclic convolution using either Rader's or Bluestein's algorithms. However, expressing the computation as a cyclic convolution requires costly steps, such as expensive memory access patterns, extra computation and/or increase in memory footprint, which impact overall performance. In this work, we present an alternative approach for such computations, where the DFT is cast as a direct symmetric matrix-vector multiplication. Using knowledge from the dense linear algebra domain and exploiting the symmetries of the Fourier matrix, we present an implementation that streams the data and requires little extra storage. We show that this approach achieves up to 2x performance gains on Intel and AMD architectures, compared to state of the art implementations offered by Intel MKL and FFTW that use Rader and Bluestein.

Keywords: Prime-Sized DFTs, Rader Algorithm, Bluestein Algorithm, Symmetric Matrix-Vector Multiplication.