Pacific Institute for the Mathematical Sciences

Hybrid Dealiasing Convolutions

Department of Mathematics and Statistical Sciences
Robert Joseph, Noel Murasko, John C. Bowman
{rjoseph1, murasko, bowman}@ualberta.ca
University of Alberta

Overview

- The convolution theorem provides a method for efficiently computing convolutions using Fast Fourier Transforms (FFTs), with the caveat that the convolution is periodic. For two arrays F and G that have period N the convolution is defined as

$$(f * g)_k = \sum_{p=0}^{N-1} f_p \odot g_{k-p}$$

- Many applications such as nonlinear PDEs require a linear convolution, which can be computed using the convolution theorem, provided that the data is sufficiently padded with zeros so we can avoid aliases.
- Implicit dealiasing provides an alternative to explicit dealiasing: the FFTs are formulated to implicitly take account of the known zero values, avoiding the need for explicit zero padding.
- To maximize performance, we use hybrid padding: the padding is chosen to be fully implicit, fully explicit, or in between.

Applications

- Convolution Neural Networks, Machine Learning, turbulence simulations, Nonlinear PDEs, signal filtering, Physics, Probability Theory usually have arrays of unequal size.

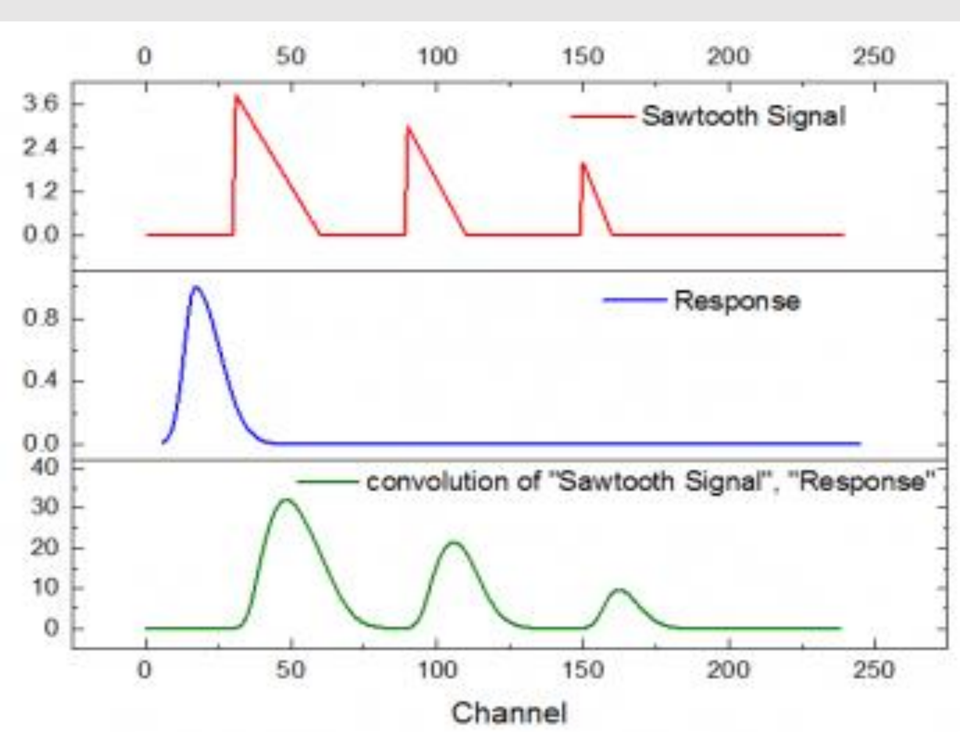

References

- Multithreaded Implicitly Dealiased Convolutions, M. Roberts and J. C. Bowman, Journal of Computational Physics 356, 98-114 (2018).

Algorithm (1 Dimensional)

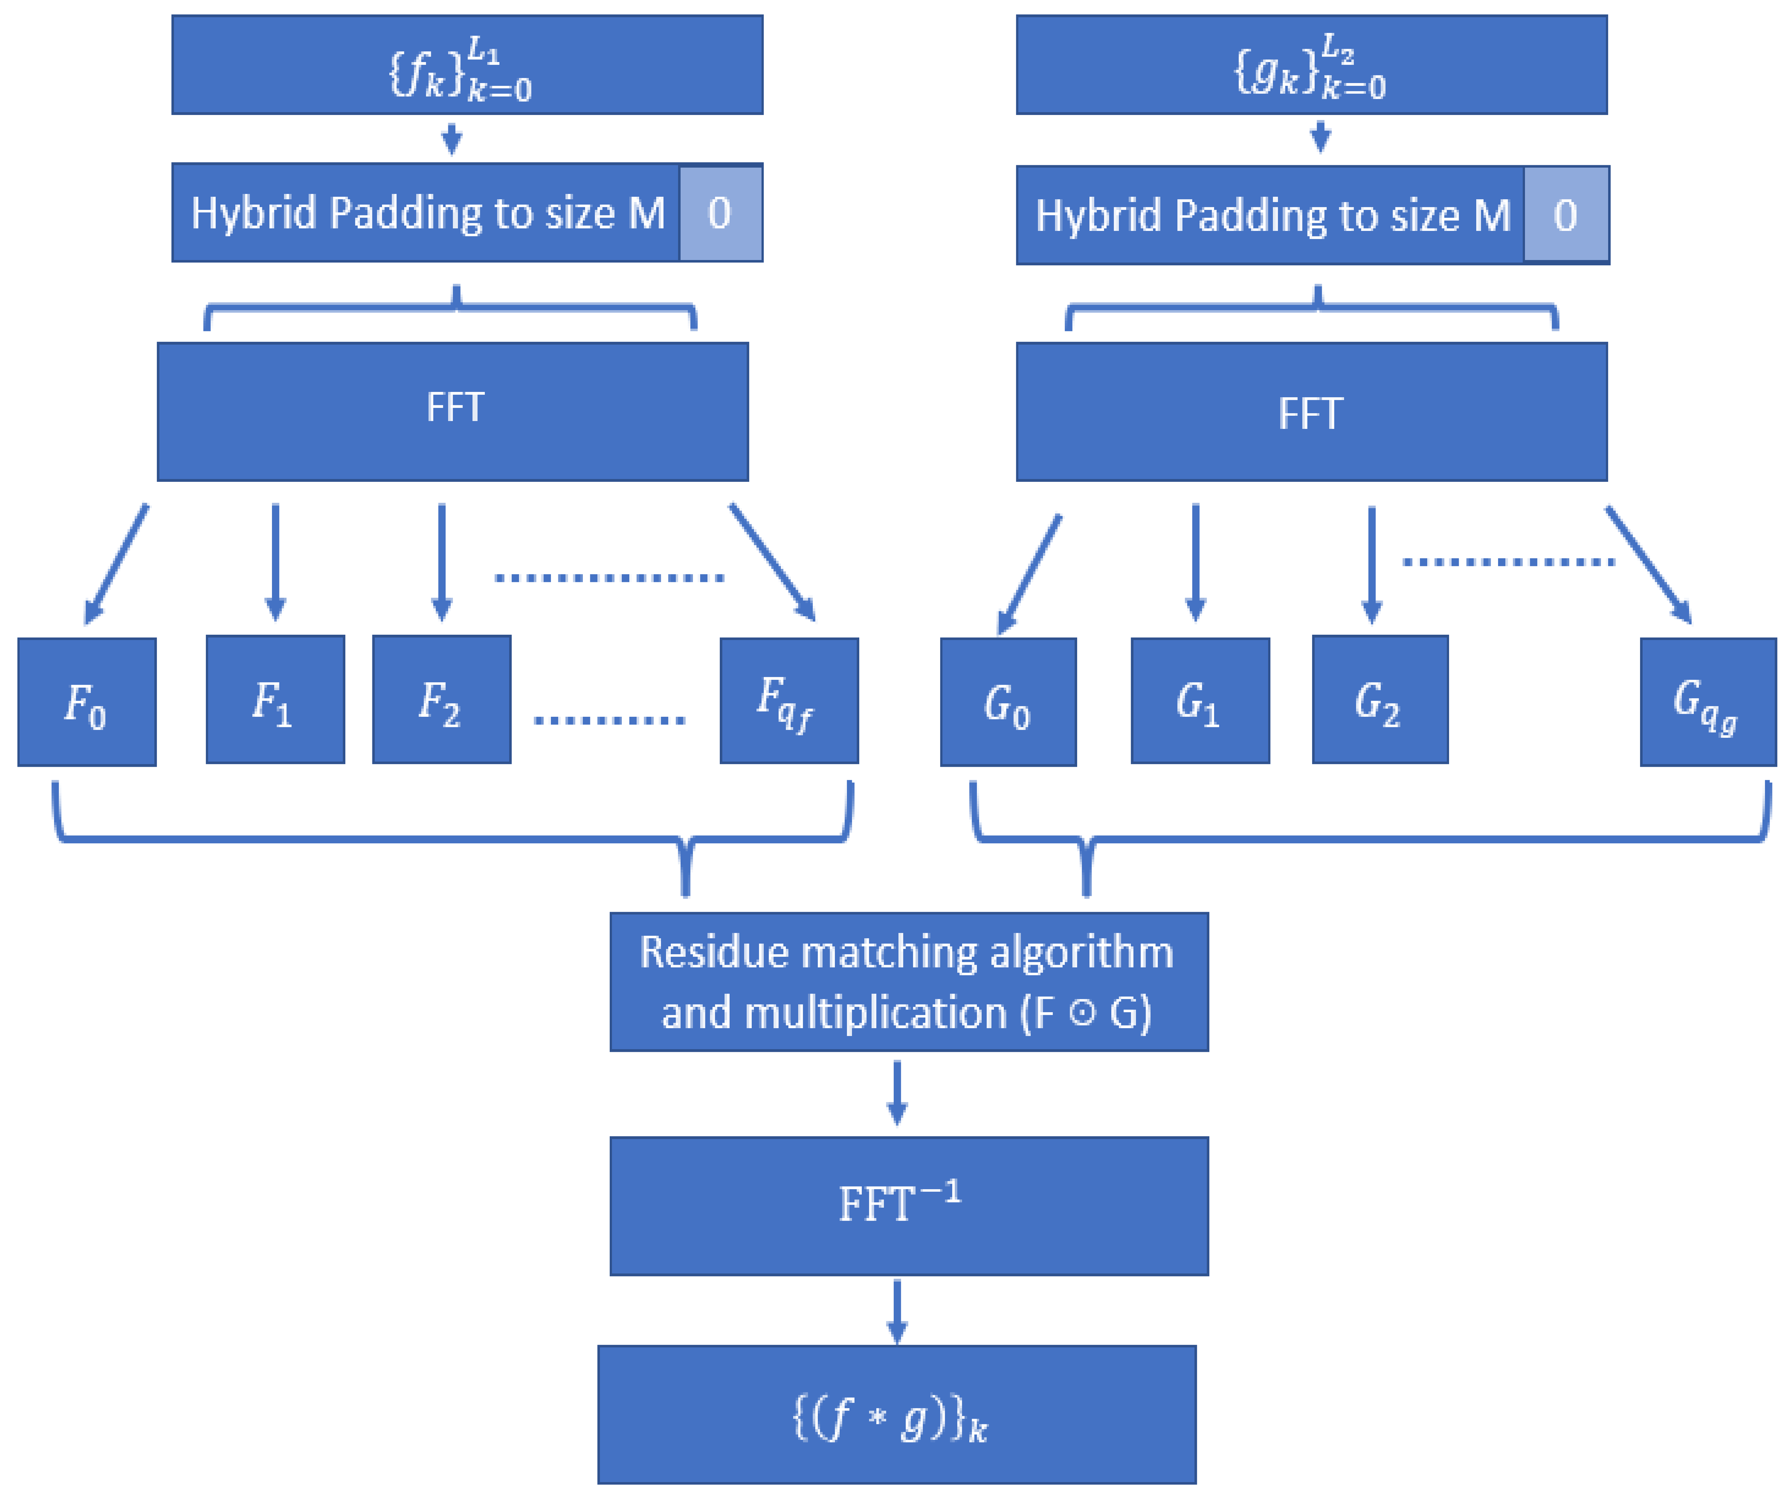

Algorithm (N Dimensional)

- For each direction we would have 3 parameters namely Λ_N, m_{f_N}, M_N where N denotes the dimension.
- Then we call the forward transform (FFT) $N-1$ dimension routine recursively along the N^{th} axis. We then multiply the two arrays in Fourier space using the Residue matching algorithm and then again call the backward transform $N-1$ dimension routing recursively along the N^{th} axis.
- Finally we use the Mode selection algorithm to return how much of the convolution we want.

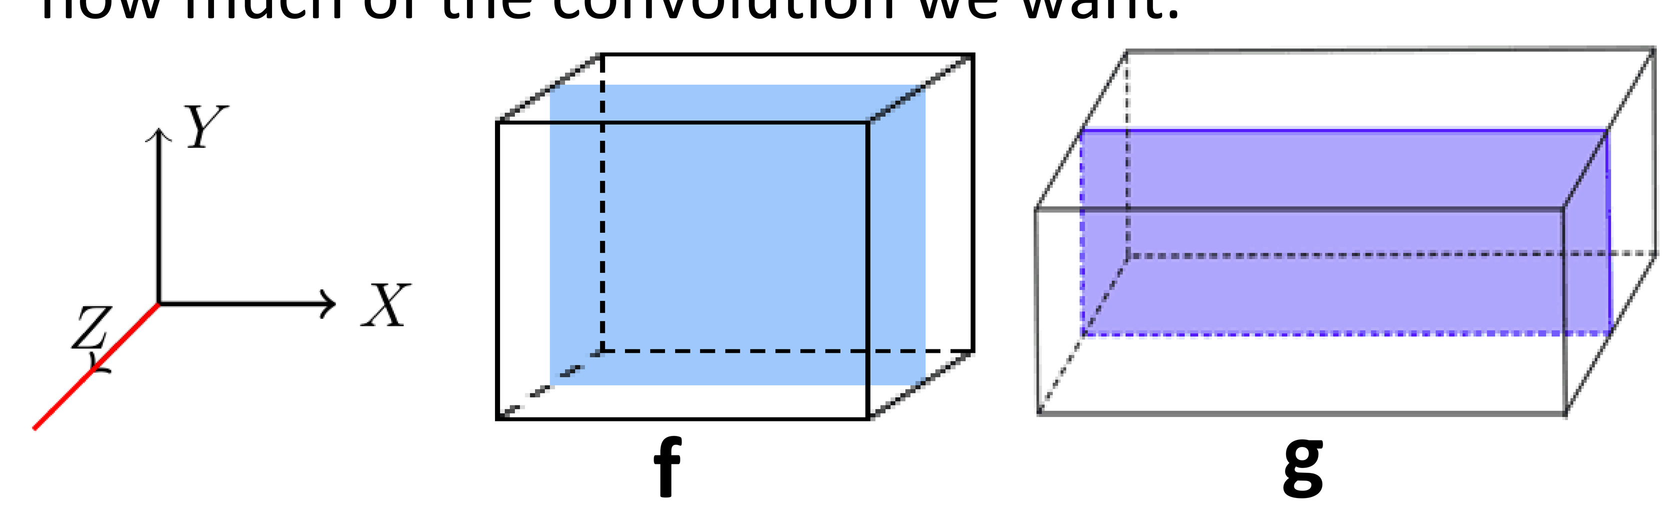

Tuning Parameters

- To find the best Λ_N, m_{f_N}, M_N along the N^{th} axis dimension we use a grid search across the factors of the length of the f and g array and hence the time complexity would be O(n) where $n = |g|$.
- We do this independently for each dimension due to the very nature of how a multi-dimensional FFT is calculated and is optimal for any dimension.

Results

- We offer unique solutions that expand hybrid dealiasing to issues with uneven input reducing memory and computation time.
- Expanded this to multi-convolution (both 1D/2D), applying the convolution to a sequence of data, N arrays rather than just two.
- We construct the first hybrid dealiasing solution, which is around 10 times quicker than typical explicit padding approaches.

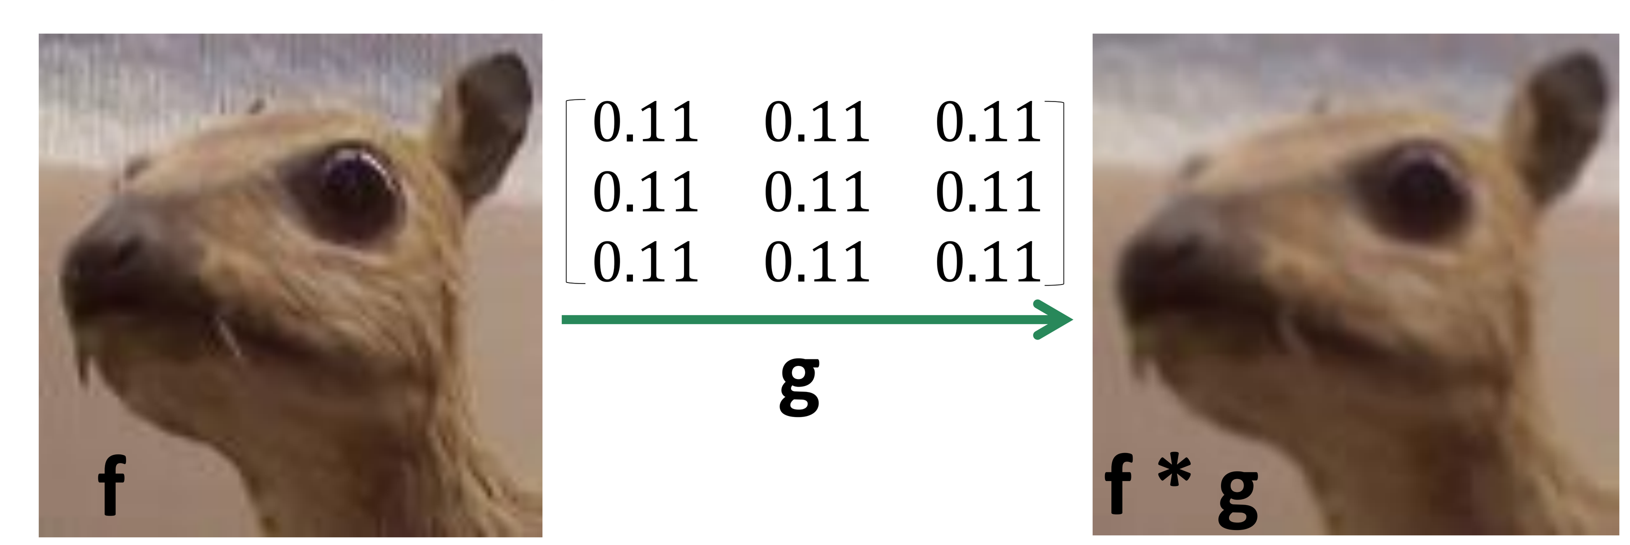

Future Work

- Complete 3D unequal hybrid dealiasing.
- Port to C++ and include in convolution libraries such as FFTW++; create wrappers for Python and Julia.
- Create specialized algorithms for the Hermitian and real case to save additional memory.
- Generalize if possible to N dimensions.
- Scan QR code for contact information.