

High Performance Medical Image Reconstruction. What could go wrong?

Abstract:

Medical Imaging devices are capable of generating many gigabytes of data per second. The best methods of reconstructing that data require sophisticated optimization algorithms and parallel computation. But unlike glitches in your favourite birds app, mistakes here could kill someone.

There are two sources of problems in such software: errors in implementation, and errors in algorithm selection or parametrization. Errors in implementation are all avoidable with the proper application of mathematical program construction and verification, but are not currently considered cost-effective, and creating interdisciplinary teams capable of mastering the required tools and techniques would be challenging. The Coconut (COde CONstructing User Tool) Project is developing integrated tools to solve this problem. Errors in algorithm selection are a more difficult to characterize, and are not even sufficiently recognized in academia, where methods with increasingly opaque failure modes, like Compressive Sensing are the focus of attention. As a first step towards avoiding problems of algorithm selection, we have developed tools for evaluating the numerical efficiency of reconstruction techniques and experiment design, starting with basic imaging experiments in Magnetic Resonance Imaging (MRI).

To eliminate implementation errors, Coconut will provide each domain expert collaborating on high-performance, safety-critical scientific software their own interface allowing for the separate specification of all ingredients from mathematical models down to efficient instruction scheduling, with all transformations providing proofs of correctness and safety. This talk will provide (1) a description of the highest level of abstraction, the Coconut Expression Library (CEL), in which domain experts can transparently specify mathematical models and regularizers, independent of implementation considerations, and of subsequent optimized code generation---including algebraic simplifications, symbolic differentiation, common subexpression elimination (CSE), and parallelization---via term graph transformation rules; (2) a sketch of our abstraction for parallel programs which affords linear-time safety proofs; and (3) a prototype type system for CEL which would allow the compiler to flag mathematical modelling errors, eliminating whole classes of errors from scientific programs.

To improve algorithm selection including the underlying experiment design, the Smallest Singular Value (SSV) metric is a computationally efficient measure of the numerical efficiency of the reconstruction method and experiment design. Although we are a long way from proving the fitness to purpose, by quantifying (and maximizing) the signal to noise ratio (SNR) in the resulting images, we are able to make the first step in this direction. To understand the importance of the SSV, we need to review basic MRI physics, to set up the resulting inverse problems, and explain how the need for faster acquisition times complicates this process, and leads to interesting questions of experiment design. Finally, we show that the SSV agrees with the gold standard for evaluating acquisition strategies, the g-factor, but with much better performance.