Project Overview

Efficient Algorithms and a Compiler for Parallel Nonlinear Stencil Computations

Many high-end scientific applications perform ***stencil computations*** in their inner loops, where a ***stencil*** is a pattern that is used to compute the value of a cell in a d-dimensional spatial grid at a timestep from the values of the nearby cells at previous timesteps. The state-of-the-art approaches to implement stencil computations are by using *cache-aware tiled algorithms* or *cache-oblivious recursive divide-and-conquer algorithms*, as these algorithms exploit better data locality without affecting the computational complexity.

Recent research by PI's group has shown that it is possible to achieve ***polynomial improvements*** in computational complexity over state-of-the-art stencil algorithms for *linear* stencils with arbitrary boundary conditions by using *fast Fourier transforms (FFT)*. This is the first time that FFT is being used to perform generic stencil computations to evolve a given spatial grid data over multiple timesteps. Very recently the same group has also shown that it is possible to compute a restricted class of linear stencils even faster through a novel use of Gaussian approximations and n-body computations.

While the stencil algorithms designed recently by the PI's group achieve significant theoretical and practical performance improvements over state-of-the-art algorithms, they have their limitations. Indeed, those results apply only to linear stencils while many stencils used in important practical applications are nonlinear.

The current project will explore ways to extend the techniques used by the PI's group to speed up linear stencil computations to computations involving stencils that are nonlinear. New techniques will be invented, if necessary. A preliminary design of a domain-specific language embedded in C++ that can express nonlinear stencil computations concisely and can be compiled automatically into highly efficient parallel algorithmic codes (initially for multicore machines) will also be made. The goal is to use the seed grant to make a few initial breakthroughs in speeding up nonlinear stencil computations and then apply for external funding to extend/generalize those initial results to solve real-world nonlinear stencil problems efficiently.

This proposal represents an original direction of theoretical and practical research in automatically generating highly efficient stencil programs.

The state-of-the-art (cache-aware tiled and cache-oblivious) stencil algorithms improve over the simple nested loop algorithms by exploiting data locality, but not affecting computational complexity. In contrast, the new algorithms will improve over the state-of-the-art stencil algorithms by lowering the computational cost asymptotically.

Broader Impacts

This research will enable scientific researchers and ordinary programmers to easily produce highly efficient codes for complex stencil computations on multiple hardware platforms using a standalone stencil compiler. A wide variety of stencil-based applications — ranging across physics, biology, chemistry, energy, climate, mechanical and electrical engineering, finance, and other areas — will become easier to develop and maintain, benefiting these application areas, as well as society at large.

Keywords: stencil computation; compiler; FFT; FMM; domain-specific compilation; parallel computing.