A SULI Summer Experience at ORNL:
Data-aware Optimization of PARADYSE

Chris Swindell

MS Candidate, TTU Electrical and Computer Engineering

Neutron scattering is a state-of-the-art technique for material structure determination, producing scattering patterns containing information unique to the internal structure of a material. Processing this data is a computationally challenging and iterative process, typically requiring between hours and weeks to complete. A key mathematical kernel that is computed during each iteration of the data analysis is referred to as the Debye summation and is well-recognized as a very computationally expensive bottleneck. To address this computational challenge, a novel parallel algorithm was developed recently to speed up the computations of Debye sums and implemented into a prototype called Parallel Debye Summation Engine (PARADYSE). My task was to integrate new methods into the PARADYSE framework to optimize its parallel execution performance and scale its applicability to material systems with larger numbers of atoms than was previously possible. PARADYSE is written in the C programming language and uses the Message Passing Interface as the primary inter-process communication handler. As part of the modifications implemented in this project, novel solutions were developed to enable data-aware parallel executions of PARADYSE with significantly reduced memory requirements, making Debye sum computations of larger atomic systems possible. The optimized version of PARADYSE was tested on a multicore Linux system with two 4-core Intel Xeon processors and on a single Summit node. Compared with a baseline sequential code that uses exact methods to compute the same scattering profiles, PARADYSE was demonstrated to deliver 98X and 734X speedup, respectively. I will discuss the testing of data-aware techniques to realize these gains and also talk about my experience in the Science Undergraduate Laboratory Internship (SULI) summer program.