

# COMPUTATIONAL RESEARCH in BOSTON and BEYOND SEMINAR

## Accelerating the Jacobi Iteration for Solving Linear Systems of Equations using Theory, Data and High Performance Computing

**MOHAMMAD SHAFIET ISLAM**

Massachusetts Institute of Technology

### ABSTRACT:

High fidelity scientific simulations modeling physical phenomena typically require solving large sparse linear systems of equations which result from the discretization of a partial differential equation (PDE) by some numerical method. The solution of these linear systems often takes a vast amount of computational time to compute. Solving these linear systems efficiently requires the use of massively parallel hardware with high computational throughput (such as GPUs), as well as the development of linear solver algorithms which respect the memory hierarchy of these hardware architectures to achieve the best performance.

This talk offers two key components towards the development of a memory efficient linear solver algorithm tailored towards high performance computing (HPC) systems. Firstly, starting with the Jacobi iteration (a parallel linear solver algorithm well-suited for HPC), we develop a family of relaxation schemes which greatly improve the convergence of the method. These schemes, termed Scheduled Relaxation Jacobi (SRJ) schemes, provide acceleration for both symmetric and nonsymmetric linear systems of equations. In the symmetric case, a data informed heuristic is developed to aid scheme selection in a practical implementation without user intervention. Secondly, we develop a high-performance GPU implementation of the Jacobi iteration method. The main characteristic of the linear solver is that it utilizes on-chip shared memory for improved memory efficiency. This is enabled by the unstructured swept rule, an algorithm for space-time decomposition which enables efficient stencil computations in parallel on unstructured grids. The shared memory Jacobi linear solver demonstrates improved performance over a classical GPU implementation which relies solely on global memory for solving two-dimensional unstructured problems. These contributions provide the basis for an efficient GPU linear solver for the solution of (potentially unstructured/nonsymmetric) linear systems arising from simulation.

**FRIDAY, MARCH 3, 2023**

**12:00 PM – 1:00 PM**

*<https://math.mit.edu/sites/crib/>*

***ZOOM Link:***

**<https://mit.zoom.us/j/96155042770>**