Orchid: Towards Heterogeneous Batched Eigenvalue Solvers

Matthew Chung
University of California, Riverside
USA

Mohammad Alaul Haque Monil (Advisor), Keita Teranishi (Advisor), Narasinga Rao Miniskar (Advisor)
Oak Ridge National Laboratory
USA

Abstract

There is a growing need for the efficient solution of many small eigenvalue problems (up to $N = 1500$) that arise in emerging scientific applications. These small-to-medium sized problems present unique computational challenges, particularly when thousands or millions of such problems must be solved repeatedly. This work presents Orchid, a novel distributed, heterogeneous, batched eigenvalue solver based on the IRIS runtime. Orchid can utilize all compute platforms in both heterogeneous nodes and clusters by harnessing the capabilities of the IRIS architecture. Orchid leverages heterogeneous architectures across multiple nodes by partitioning the application task DAG intelligently and orchestrates multiple instances of the IRIS runtime via MPI. We evaluate our proposal against two heterogeneous hardware configurations and Frontier, demonstrating Orchid's performance utilizing both intra-node and inter-node heterogeneity.

1 Introduction

The symmetric eigenvalue problem is fundamental in computational science, as many physical phenomena can be expressed in terms of real modes and energies. Because symmetric matrices have real spectra and orthogonal eigenvectors, these problems admit stable algorithms and physically meaningful interpretations.

For example, in quantum mechanics, the stationary Schrödinger equation reduces to the Hermitian problem

$$H\psi = E\psi,$$

where eigenvalues E represent energy levels and eigenvectors ψ describe electronic states. In structural mechanics and seismology, the generalized eigenvalue problem

$$Ku = \lambda Mu,$$

with stiffness K and mass M, yields natural frequencies $\omega = \sqrt{\lambda}$ and vibration modes u.

Traditionally, research has emphasized solving very large eigenvalue problems at extreme scales. However, there is also a growing need for the efficient solution of many small eigenvalue problems (up to $N = 1500$) that arise in emerging scientific applications. These small-to-medium sized problems present unique computational challenges, particularly when thousands or millions of such problems must be solved repeatedly.

The main motivation and objective of this work is to provide a distributed batched eigenvalue solver for extremely heterogeneous compute environments. This framework looks to maximize productivity and computational efficiency regardless of the underlying hardware. There are two major challenges that are addressed in this work: integrating the eigenvalue solver into the IRIS programming model in a productive manner and extending the existing IRIS [3] runtime to performantly support distributed computation.

2 Approach

2.1 MatRIS Integration

Integration of the EigenG-Batched [1] solver into IRIS is accomplished through adding a custom batched eigenvalue solver to the MatRIS [4] library. MatRIS is a high performance and portable heterogeneous math library designed specifically for the IRIS runtime. MatRIS provides a collection of portable BLAS and LAPACK routines in the form of kernels capable of being executed on heterogeneous platforms. We extend MatRIS to include a tiled batched eigenvalue solver capable of execution in a multi-node heterogeneous environment. The eigenvalue solver splits the batched workload into L tiles where each tile is represented as an IRIS task. This approach enables parallel execution and reduces the memory requirement for each task. We also provide CUDA and HIP kernels, both utilizing the EigenG-Batched algorithm, allowing for IRIS to simultaneously execute the solver on both GPU platforms.

2.2 Global Coordinator

We add a lightweight MPI-based global coordinator that manages the high level control flow. The global coordinator is responsible for partitioning the application DAG on rank 0, initializing the IRIS runtimes on other nodes, sending each node its DAG partition, and collecting/aggregating results at the end.

2.3 DAG Partitioning Algorithm

Partitioning the IRIS task DAG across heterogeneous nodes presents a unique challenge. The partitioning process must be capability aware, or risk catastrophic performance loss due to computational bottlenecks. We propose using a customizable multi-constraint graph partitioning algorithm implemented using METIS [2]. METIS is a graph partitioning algorithm whose objective is to minimize total cut weight while conforming to n constraints. METIS allows developers to provide hints regarding data locality between dependent tasks as well as per-task computational requirements by adjusting edge weights and vertex weights respectively. METIS also allows IRIS to be capability aware across multiple constraints (GPU/CPU count, GPU Memory, TFLOPs) by adjusting the target weights. With this metadata, METIS partitions k-ways (k = number of nodes), resulting in partitions that are sized proportionally to the capabilities of each node while also minimizing expensive cross-node dependencies. With METIS, IRIS can maximize device utilization and reduce communication overhead throughout the application makespan.

3 Evaluation

To evaluate our batched eigenvalue solver's functionality in a intra-node heterogeneous environment, we run a workload of 4096 symmetric matrices of size 256 against a node with AMD CPU, a NVIDIA RTX 3090, and a AMD Radeon RT 6900XT. With only the MatRIS host code, both GPUs execute their batched kernels concurrently, with high utilization across both devices. We also test against a workload of 8192 matrices of size 64, which also shows effective scalability in a node with multiple architectures. To demonstrate Orchid's computational performance, we evaluate performance of the EigenG-Batched against the reference batched eigenvalue solvers provided by cuSOLVER and hipSOLVER. We evaluate our performance against two platforms, a node with 2x NVIDIA H100 NVL and a node with 4x AMD MI300A. We test two workloads, 4096 matrices of size 512 and 65536 matrices of size 64. On both platforms and workloads, Orchid provides better scaling across multiple GPUs in comparison to vendor solvers. In terms of time-to-solution Orchid offers competitive or better performance. We observe certain cases where hipSOLVER fails to scale effectively across multiple GPUs, instead increasing time-to-solution when utilizing more GPUs. With EigenG-Batched's competitive performance to vendor libraries and Orchid's ability to leverage heterogeneity, we are able to provide a productive batched eigenvalue solver.

4 Summary

Symmetric eigenvalue solvers are a key part of computational science because of their role in modeling physical phenomena in large-scale simulations. However, because of their high computational cost, designing a parallel eigenvalue solver is critical to realizing the full potential of HPC resources in scientific computing. We investigate possible solutions, presenting a batched eigenvalue solver integrated with the IRIS runtime system, fully exploiting distributed and heterogeneous systems.

References

[1] Yuta Hasegawa, Toshiyuki Imamura, Takuya Ina, Naoyuki Onodera, Yuuichi Asahi, and Yasuhiro Idomura. 2022. GPU Optimization of Lattice Boltzmann Method with Local Ensemble Transform Kalman Filter. In *2022 IEEE/ACM Workshop on Latest Advances in Scalable Algorithms for Large-Scale Heterogeneous Systems (ScalAH)*. 10–17. https://doi.org/10.1109/ScalAH56622.2022.00007

[2] George Karypis and Vipin Kumar. 1997. METIS—A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes and Computing Fill-Reducing Ordering of Sparse Matrices. (01 1997).

[3] Jungwon Kim, Seyong Lee, Beau Johnston, and Jeffrey S. Vetter. 2021. IRIS: A Portable Runtime System Exploiting Multiple Heterogeneous Programming Systems. In *Proceedings of the 25th IEEE High Performance Extreme Computing Conference (HPEC "21)*. 1–8. https://doi.org/10.1109/HPEC49654.2021.9622873

[4] Mohammad Alaul Haque Monil, Narasinga Rao Miniskar, Keita Teranishi, Jeffrey S. Vetter, and Pedro Valero-Lara. 2023. MatRIS: Multi-level Math Library Abstraction for Heterogeneity and Performance Portability using IRIS Runtime. In *Proceedings of the SC '23 Workshops of the International Conference on High Performance Computing, Network, Storage, and Analysis* (Denver, CO, USA) *(SC-W '23)*. Association for Computing Machinery, New York, NY, USA, 1081–1092. https://doi.org/10.1145/3624062.3624184