Grace Hopper Superchip

Table of contents

Parabricks is proud to announce full optimization and support for the groundbreaking Nvidia Grace Hopper superchip. The Nvidia GH200 Grace Hopper Superchip combines the Nvidia Grace and Hopper architectures using Nvidia NVLink-C2C to deliver a CPU+GPU coherent memory model for accelerated Artificial Intelligence (AI) and High Performance Computing (HPC) applications. This integration represents a significant leap forward in genomic data analysis, allowing researchers to tackle complex analyses with unprecedented speed and efficiency. The Nvidia Grace Hopper Superchip is the first true heterogeneous accelerated platform for HPC and AI workloads. It accelerates applications with the strengths of both GPUs and CPUs while providing the simplest and most productive programming model for performance, portability, and productivity.

Key Features of the Nvidia Grace Hopper Superchip GH200

Feature	Description
Grace CPU cores (number)	Up to 72 cores
CPU LPDDR5X bandwidth (GB/s)	Up to 500GB/s
GPU HBM bandwidth (GB/s)	4TB/s HBM3, 4.9TB/s HBM3e
NVLink-C2C bandwidth (GB/s)	900GB/s total, 450GB/s per direction
CPU LPDDR5X capacity (GB)	Up to 480GB
GPU HBM capacity (GB)	96GB HBM3, 144GB HBM3e
PCIe Gen 5 Lanes	64x

By harnessing the immense computational capabilities of the Nvidia Grace Hopper Superchip, users can experience even greater acceleration and throughput for their genomic pipelines.

Documentation

All tools and pipelines from Parabricks 4.3.1-1 are now optimized and supported on the Nvidia Grace Hopper Superchip, therefore we refer the users and developers to the Tool Reference.

Performance tuning

To achieve optimal performance for all Parabricks tools on the Nvidia Grace CPU we refer the users and developers to the Grace CPU benchmarking guide. This guide will illustrate recommendations and best practices directly related to the Nvidia Grace CPU and help you realize the best possible performance for your particular system.

Get Started

- To begin using Parabricks and unleash the power of the Nvidia Grace Hopper Superchip for your genomic analyses, you can obtain the docker image running the following command:

  ```
  $ docker pull nvcr.io/nvidia/clara/clara-parabricks:4.3.1-1.grace
  ```

  and follow the Tutorials.

- For any questions or support, please visit the Nvidia Parabricks Community. Join a vibrant community of researchers and experts to exchange ideas, seek assistance, and stay updated on the latest developments in genomic data analysis.

- To learn more about the Nvidia Grace Hopper Superchip please visit here.