Zhiyu Ding

dingzhiyu2004@163.com · (+86) 18765788600 · nevergpdzy.com

Education

Southwest Petroleum University, Chengdu, Sichuan September 2023 – Present

Undergraduate Student in Data Science and Big Data Technology, School of Computer and Software Engineering

- **Academic Performance**: GPA: 4.13/5.0, Major Ranking: 1/66
- **Languages**: English CET-6 (478), CET-4 (521)
- **Core Courses**: Advanced Mathematics II(95), Linear Algebra(91), Probability and Statistics(93), Principles of Statistics(93), Data Structures and Algorithms(90)

Project Experience

Parallel Computing Optimization for Oil Spill Prediction Model

May 2024 – August 2024

2024 Marine Computing Challenge Finals Team Leader

Selected for the 2024 Marine Computing Challenge finals; led hybrid parallelization and coastline collision detection optimization for a 2D oil spill prediction model in a dual-node 128-core environment.

- Completed MPI+OpenMP hybrid parallelization on the original serial program, analyzing bottlenecks with VTune and optimizing via load balancing, non-blocking communication, packed communication, disjoint segment rapid elimination, and binary search
- Passed the organizer's correctness verification, achieved a 2482.14x speedup over the baseline, ranked 5th in the finals, and won the national third prize

Tecorigin Deep Learning Operator Performance Optimization

June 2024 – December 2024

2nd OpenAtom Competition – Tecorigin Operator Development Challenge Team Leader

Selected for the 2nd OpenAtom Competition Tecorigin Operator Development Challenge; conducted performance optimization of Tecorigin convolution forward operators on TaiChu accelerator.

- Built performance profiles using Perf Data, identifying I/O as the core bottleneck at 93.1%; designed a 234KB SPM management strategy and double-buffering asynchronous pipeline optimizing memory access, and implemented SIMD data reordering based on floatv16 vector registers, significantly improving write-back bandwidth
- Reduced total execution time from 1820.78ms to 489.18ms, achieving a 3.7x overall speedup; resolved the mismatch between matrix multiplication library output and NHWC format, winning the national third prize

PCG Algorithm Optimization on New Generation Sunway Supercomputer

February 2024 – April 2024

7th Domestic CPU Parallel Application Challenge Individual Project

Selected for the 7th Domestic CPU Parallel Application Challenge preliminary round; optimized the Preconditioned Conjugate Gradient (PCG) solver for the Sunway heterogeneous many-core architecture.

- Restructured the PCG workflow based on the Sunway 64-core slave-core architecture, parallelizing core operators such as SpMV, dot product, and preconditioning using athread, and optimizing data access and blocking strategies via DMA and on-chip LDM
- Through over 40 iterations, reduced runtime from 1287 seconds to 32.5 seconds, achieving a 39.6x speedup

Research Experience

Image Processing and Parallel Computing Laboratory (IPPC Lab), Southwest Petroleum University

May 2024 – September 2025

3D Ultrasound Elastography CUDA Acceleration Sichuan Student Innovation Program; Advisor: Prof. Bo Peng

- Rebuilt the serial MATLAB workflow as a CUDA C++ parallel pipeline for high-volume cross-correlation in 3D ultrasound elastography, decomposing it into RF integral volume construction, cross-correlation integral volume generation, and 3D subpixel displacement estimation GPU stages; successfully ported the algorithm to Sugon DCU accelerator
- Optimized memory access for large-scale 3D prefix-sum computation and introduced warp-level cooperation, accelerating the core kernel by 11.11x and reducing end-to-end runtime from 16.32s to 0.22s, about 73x faster

CUDA Parallel Optimization for Quadtree Adaptive Image Tiling Algorithm

October 2025 – March 2026

- Addressed computational bottlenecks in large-image adaptive tiling for Visual Transformer pretraining by fully parallelizing the quadtree adaptive tiling pipeline in CUDA C++, including integral image construction, quadtree iterative merging, hierarchical noise injection, and patch synthesis; employed Top-K selection to process top-k candidates per iteration, avoiding full sorting
- Decomposed the pipeline into integral image computation, quadtree partitioning, hierarchical noise generation, and patch synthesis, processing 8192-pixel edge-length images on A100 in 3.16s with 13.0x speedup over serial CPU version, and 69.1x faster noise generation

Awards and Honors

International Second Prize, ASC2025 World Student Supercomputer Challenge June 2025
Invited Participant, SC24 International Supercomputing Competition Online Track IndySCC November 2024
National Third Prize, 2024 Marine Computing Challenge Finals August 2024
National Third Prize, Tecorigin Operator Development Challenge Finals December 2024
Provincial Second Place, Tianyi Cloud Xirang Cup College AI Competition, Sichuan Provincial Competition July 2025
National Third Prize, 15th Blue Bridge Cup National Finals June 2024
First-class Scholarship, Second-class Scholarship, Southwest Petroleum University Outstanding Student Scholarship 2024

Additional Information

- Blog: https://nevergpdzy.com
- GitHub: https://github.com/NeverGpDzy
- Research Interests: High Performance Computing and Parallel Programming, GPU/CPU heterogeneous computing optimization