GPU Programming with OpenMP

Dr. Hermes Senger **Dr. Roussian R.A. Gaioso**

email: hermes@ufscar.br
Departmento de Computação
Universidade Federal de São Carlos- UFSCar

INFIERI 2023
August 28th to September 9th 2023
São Paulo – SP, Brazil

Why program GPUs?

The Pros

- Thousands of parallel floating-point units for a few hundred dollars $
- Outstanding performance & energy efficiency for many applications (e.g., numerical applications, multimedia, machine learning & AI, etc)
- Can be plugged into PCs to Supercomputers (seven out of the ten top clusters of the www.top500.org derive the lion's share of their compute power from GPUs)

And (a few) cons

- Applications must be designed/optimized for GPUs
- Programming for GPUs is (a little bit) harder to program than CPUs

How to program GPUs?

- Many programming languages, libraries, etc: CUDA, OpenAcc, OpenMP, SYCL, HIP, pyCUDA, cuPy, Numba, etc
- OpenMP
  - C/C++ and FORTRAN interfaces
  - Directive-based approach: **#pragma omp target** construct to offload code for GPUs
  - Good availability of compilers for CPUs and GPUs: AMD, ARM, Barcelona, Flang, Fujitsu, GNU, HPE, IBM, Intel, LLVM, Microsoft, NAG, Oracle, PGI, Siemens, PARC …

Course Overview

- OpenMP overview
- Early steps with OpenMP
- The host/device model
- Moving data to/from the device
- Supporting massive parallelism
- Code optimization exercises

Hands-On

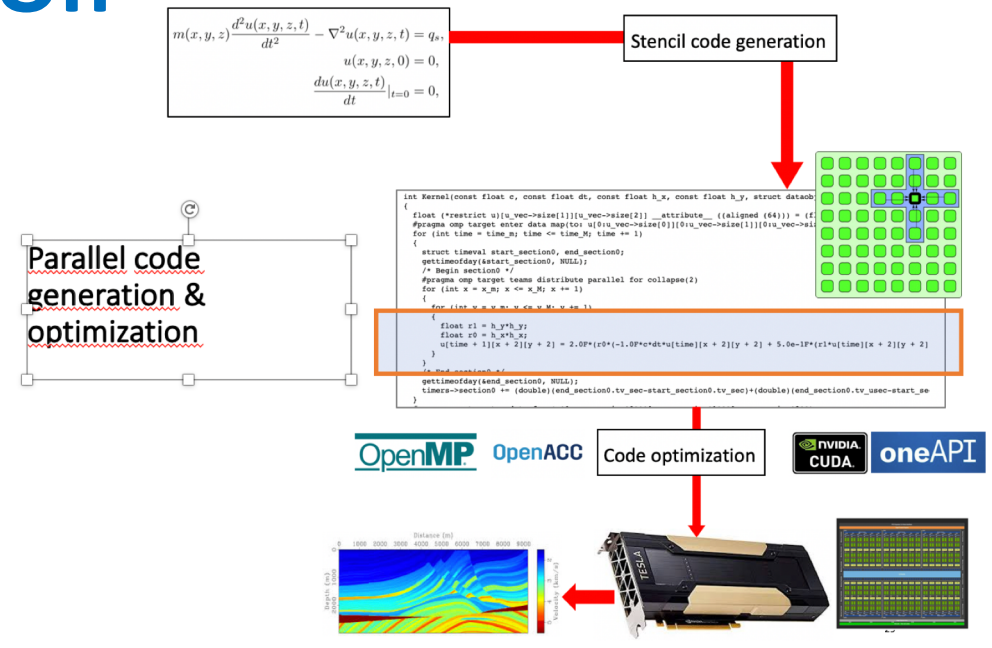

Hands-on approach

- Take the sequential code
- Parallelize/optimize
- Test the performance

Requirements:

- Basic programming skills
- A personal computer or laptop with a web browser and access to the internet
- A Google login will be necessary for students to access GPUs from the **colab** environment for the hands-on sessions

GPU Programming with OpenMP

Dr. Hermes Senger **Dr. Roussian R.A. Gaioso**

email: hermes@ufscar.br
Departmento de Computação
Universidade Federal de São Carlos- UFSCar

INFIERI 2023
August 28th to September 9th 2023
São Paulo – SP, Brazil