

[Registration link](#)

Date: November 30th to December 2nd, 2020

What is GPU Hackathon

As part of the workshop every year, KSL facilitates a coding event in which teams of developers port their own applications to run on GPUs, or optimize their existing GPU applications. Each team usually consists of 3 to 5 developers who are intimately familiar with the application and workflow. They work with one or more GPU programming experts mentoring these sessions. Our mentors come from NVIDIA and KSL supercomputing applications team.

The hackathon will span over 2 days, with 3rd day dedicated to reporting and prize ceremony. There are a variety of “suitable” codes to port to GPUs. They range from calling GPU enabled libraries in your code to refactoring using a parallel programming model for GPUs, e.g. introducing OpenACC, OpenMP for offloading or even writing a CUDA version. Parts of pythonic workloads conducive to run on GPUs can make use of packages like Numba and CuPy.

Why participate

- Explore if your code can run or optimize on GPUs to achieve faster time to solution
- Allow yourself to set aside time for training and coding sprint, in the presence of experts for face to face consulting on your code
- By the end of the hackathon, each team will ideally have ported some part of their code to run on GPUs or have a clear roadmap of how to/what to do next.

How to apply/participate

You will need to write a short description to introduce your application and project. This may include:

- Background: about your team, your application/code, science/industry problem you are addressing, etc.
- A brief status of your code: programming language, scalability, HPC use vs cloud, GPU use etc.
- What is the motivation to speeding up your code? How will it impact your project’s goals?
- Submit your proposal

Please fill the [registration form](#) to participate.

How to Prepare:

Given your proposal is accepted, you will be invited to share code, preferably as a repository hosted in e.g. GitHub, GitLab, Bitbucket, SVN etc.

If your code is more than 1000 lines, please pick a few representative kernels to start working on during the event. You may want to include where the code spends most of its compute time.

Test a representative case with valid answers before coming to the event. It is important because the workflow will include iterations of optimizations and testing both for validity and for performance.

Send your application to training@kaust.edu.sa

Deadlines:

Submission: Nov 24th, 2020

Invitation to attend: Nov 28th, 2020

Mentor must be able to access to code repository: Nov 24th, 2020