

TUTORIAL PROPOSAL

- Name of the Tutorial

Accelerating Computing Using SYCL Programming for GPUs

- Name, country, and affiliation of instructor

CLARA ALMEIDA, UC Santa Cruz, CA, USA

DAVID GUERRERO, UC Santa Cruz, CA, USA

RAYMUNDO HERNANDEZ-ESPARZA, Argonne National Laboratory, IL, USA.

MARIA PANTOJA, California Polytechnic State University, CA, USA

ERIK PAUTSCH, Loyola University Chicago, IL, USA


SILVIO RIZZI, Argonne National Laboratory, IL, USA

GEORGE K. THIRUVATHUKAL, Loyola University Chicago & Argonne National Laboratory, IL, USA

ALVARO VAZQUEZ-MAYAGOITIA, Argonne National Laboratory, IL , USA.

JORGE VELEZ, University of Illinois Urbana-Champaign, IL , USA.

- A recent picture of Instructor (for the website, at least 200x200 pixels)



- Short description of the tutorial (max 200 words)

In the last decade, the adoption of graphic processing units (GPUs) in high-performance computing has facilitated significant advancements across multiple scientific domains. The impact is evident, from enabling the development of robust large language models to powering massive simulations in biological systems.

With the emergence of new device vendors in the GPU market, such as Nvidia, AMD, and Intel, users now have expanded options in terms of cost efficiency and performance for their specific applications. Despite this, the computational science community faces challenges in adopting new architectures and writing performant and portable code to utilize these advancements effectively.

In this workshop, we will provide a concise introduction to SYCL, a parallel programming model designed to develop portable C++ code. With SYCL, developers can efficiently offload computations to accelerators, particularly GPUs. We will cover the fundamental concepts of GPU development, including a comprehensive walkthrough of exemplary cases. Additionally, we will discuss the limitations and strategies for migrating existing GPU codes to SYCL. Prior familiarity with C++ and command-line editors is expected for students attending this workshop.

- Duration 4 hours

Brief organization of tutorial

- Introduction . 1 hour. (using slides)
 - Heterogeneous computing and SYCL
- Hands-on Tutorial 55 min.
 - Evaluation of Electron Density in Real Space Grid with SYCL
- Hands-on Tutorial 55 min.
 - Porting the Marching Cubes visualization algorithm from CUDA to SYCL
- Hands-on Tutorial 55 min.
 - Implementing an N-body simulation in SYCL
- Wrap up and closing remarks 15 min
- Student's prerequisites (for instance, "High level of Python" or " Linux Administrator")

Audience:

The intended audience is intermediate students in computer science with knowledge of C/C++ programming. We can do it in English or Spanish, depending on the preference of the audience.

Students should create an account here: <https://console.cloud.intel.com/>

- Special conditions for accessing the tutorial. For instance: "Laptop with X Linux installed" or "registration for X website previous to the tutorial". Remember we don't have rooms equipped with computers, just Internet Access. Students must use their personal equipment

Attendees will require a laptop computer that can view and run jupyter notebook specifically `\url{https://console.cloud.intel.com/}`.

Links: <https://unoapi.org/index.html>

<https://www.khronos.org/sycl/>

<https://github.com/codeplaysoftware/syclacademy>