

Workshop: School of Data Science and Computation Thinking

Introduction to Parallel Computation with GPUs: Programming with JAX, NUMBA, THURST, and CUDA

While Moore's law continues to remain valid, processor speed has reached a plateau for quite some time owing to the limitations of CMOS technology. High performance computing increasingly relies on massive parallelism in order to make computationally tractable bigger and bigger problems. Parallel computing has a long history, but until recently parallel machines have been the domain of a small number of well-funded supercomputing centers. Recently however, inexpensive commodity GPUs containing thousands of processors have become available.

This workshop will explore various ways to take advantage of GPUs with an emphasis on practical hands-on programming. For some applications, is it possible to use GPUs using Python libraries such as JAX. Generally, when libraries exist for specific critical tasks that need to be speeded up, it is better to use these libraries to avoid re-inventing the wheel. On the other hand, for other applications it is necessary to program the GPU directly in order to achieve significant speed-up. NVIDIA has developed a general-purpose programming language for GPUs known as CUDA that allows a broad range of problems to be run on GPUs.

This workshop will provide an introduction to parallel programming concepts emphasizing hands-on GPU programming experience. Libraries and interfaces to GPU will also be surveyed and general issues in parallel programming such as race conditions, deadlocks, livelocks, etc. will be treated. Performance tuning will also be discussed.

This online workshop will consist of 8 one-hour sessions supplemented by practical lab exercises for participants to complete on their own time. Participants should expect to devote approximately the same amount of time to the practical exercises as to the online lectures. Participants will be provided with an account on a CUDA GPU machine using Linux for the practical exercises and for experimentation. Some experience in basic Python and C programming will be assumed.

The online sessions will take place during the following days from 4-5pm: Tu 18 Feb, Th 20 Feb, Tu 25 Feb, Th 27 Feb, Tu 4 March, Th 6 March, Tu 11 March, Th 13 March. Please register using the following online form:

https://forms.office.com/Pages/ResponsePage.aspx?id=Azv6pjwKWEKEM6Eg3_zTSMxCzkIR3YpGv05EYJ70OINUQkZIUUxKNE9HV0xIUE9SSERZT0FKS1hQNy4u

Contact (for inquiries regarding the course) : bucher@sun.ac.za

Prof Martin Bucher is Directeur de recherche au CNRS in Paris, France, based at Université Paris Cité. He is also Fractional Professor at the School of Data Science and Computational Thinking at Stellenbosch. Bucher is a NITheCS Associate and an elected member of the Academy of Science of South Africa (ASSAf). He has made numerous contributions to theoretical particle physics and cosmology and to observational cosmology. He shared the 2018 Gruber Prize in Cosmology as part of the Planck Team.