

Closing session



thematic
CERN
School of Computing

14-18 June 2021
Online

**Scientific Software for
Heterogeneous Architectures**

<http://cern.ch/csc>



<http://indico.cern.ch/e/tCSC-spring-2021>

Sebastian Lopienski
CERN School of Computing director

Timetable

<https://indico.cern.ch/event/1017080/timetable/>

| Monday, 14 June 2021 | | Tuesday, 15 June 2021 | | Wednesday, 16 June 2021 | | Thursday, 17 June 2021 | | Friday, 18 June 2021 | |
|----------------------|---|-----------------------|--|-------------------------|--|------------------------|---|----------------------|--|
| 09:00 | Opening Session | 09:00 | Writing parallel software - Danilo Piparo (CERN) | 09:00 | Hardware evolution and heterogeneity - Andrzej Nowak | 09:00 | Data-oriented design - Andrzej Nowak | 09:00 | Group assignment for Track 3: Programming for heterogeneous architectures |
| 09:30 | Preparing for the HL-LHC computational challenge - Danilo Piparo (CERN) | 10:00 | Modern programming languages for HEP - Sebastien Ponce (CERN) | 10:00 | Practical vectorization - Sebastien Ponce (CERN) | 10:00 | Programming for GPUs - Dorothea vom Bruch (CPPM/CNRS) | 10:30 | Coffee break |
| 10:30 | Coffee break | 11:00 | Coffee break | 11:00 | Coffee break | 11:00 | Coffee break | 10:50 | Design patterns and best practices - Daniel Campora (University of Maastricht) |
| 10:50 | Introduction to efficient computing - Andrzej Nowak | 11:20 | Optimizing existing large codebase - Sebastien Ponce (CERN) | 11:20 | Scientific computing on heterogeneous architectures - Dorothea vom Bruch (CPPM/CNRS) | 11:20 | Performant programming for GPUs - Daniel Campora (University of Maastricht) | 11:50 | Programming for heterogeneous architectures |
| 11:50 | Self-presentation: 1 minute per p... | | | | | | | | |
| 14:00 | Self-presentation: 1 minute per p... | 14:00 | School photo | 14:00 | Group assignment for Track 2: Parallel and Optimised Scientific Software | 14:00 | Programming for heterogeneous architectures - exercise - Dorothea vom Bruch (CPPM/CNRS) Daniel Campora (University of Maastricht) | 14:00 | Summary and future technologies overview - Andrzej Nowak |
| 14:30 | Coffee break | 14:05 | Parallel and optimised scientific ... | 15:30 | Coffee break | 14:20 | Programming for heterogeneous architectures - exercise - Dorothea vom Bruch (CPPM/CNRS) Daniel Campora (University of Maastricht) | 15:00 | Coffee break |
| 14:50 | Group assignment for Track 1: Technologies and Platforms | 14:25 | Parallel and optimised scientific software - exercise - Sebastien Ponce (CERN) Arthur Hennequin (CNRS) | 15:50 | Student lightning talks | | | 15:30 | Exam |
| | | | | 16:30 | Parallel and optimised scientific s... | | | | |
| | | | | | | | | 16:40 | Closing Session |

Academic programme

<https://indico.cern.ch/event/1017080/program>

Theme: *“Scientific Software for Heterogeneous Architectures”*

- **Introduction lecture** *“Preparing for the HL-LHC computational challenge”* by **Danilo Piparo** (CERN)
- **Track 1: Technologies and Platforms** by **Andrzej Nowak**
- **Track 2: Parallel and Optimised Scientific Software** by **Danilo Piparo** (CERN) and **Sebastien Ponce** (CERN) exercises assisted by **Arthur Hennequin** (CNRS)
- **Track 3: Programming for Heterogeneous Architectures** by **Dorothea vom Bruch** (CPPM/CNRS) and **Daniel Campora** (University of Maastricht)



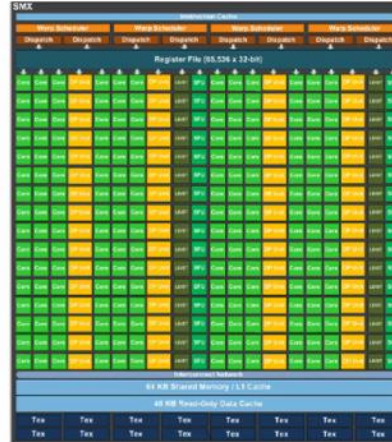
Lectures

Hardware Evolution and Heterogeneity



The NVIDIA ecosystem – hardware

- High performing range of vector devices with an acceptable cost
- Hardware:
 - Limited number of control threads – typically 16-32/unit
 - High number of parallel data elements – hundreds in total
 - Very limited control flow capabilities – adapted shaders serve as vector units, low cache
- Performance: >10 TF SP
- Tesla line with ECC support
- Double precision is lower: e.g. 1/2 of peak BW
- Full IEEE compliance available at a cost



SIMX: 192 single-precision CUDA cores, 64 double-precision units, 52 special function units (SFUs), and 32 load/store units (LD/ST).

Image: NVIDIA



Andrzej NOWAK (speaker)



Florian BURY



Sylvain FARGI...



Saswat MISHRA



Patrick CONNOR



Moises LEON...



Matteo BARBETTI

Recording

Matrix-matrix multiplication

Suppose we want to multiply two arrays:

- ▶ **A** of size $m \times n$
- ▶ **B** of size $n \times k$
- ▶ Result is **C** of size $m \times k$

```

1  __device__ void multiply_arrays(float* A, float* B, float* C, int m, int n, int k) {
2  for (int row = threadIdx.x; row < m; row += blockDim.x) {
3      for (int col = 0; col < k; ++col) {
4          float element = 0.f;
5          for (int i = 0; i < n; ++i) {
6              element += A[row * m + i] * B[i * k + col];
7          }
8          C[row * k + col] = element;
9      }
10 }
11 }

```



Lecture quizzes

What are the changes between Python 2 and Python 3? *

- Small syntactic improvements, everything is backward compatible anyway
- Major changes are present, backward incompatible, migration is almost impossible
- Major changes are present, backward incompatible, but migration is possible
- These are independent languages, both will be maintained in the future




What are the changes between Python 2 and Python 3? *

- Small syntactic improvements, everything is backward compatible anyway
- Major changes are present, backward incompatible, migration is almost impossible
- Major changes are present, backward incompatible, but migration is possible**
It is even possible to concurrently support both languages in the same code
- These are independent languages, both will be maintained in the future

Student lightning talks – tCSC spring 2021

- *"FRED: a fast Monte Carlo code on GPU for Treatment Planning Software"*
by **Micol De Simoni** (Sapienza University of Rome)
- *"Track reconstruction on heterogeneous architectures with SYCL"*
by **Bartosz Sobol** (Jagiellonian University Krakow)
- *"Matrix Element Regression with Deep Neural Networks - breaking the CF"*
by **Florian Bury** (Catholic University of Louvain)
- *"Exploring Heterogeneous Architectures in Track Reconstruction S"*
by **Georgiana Mania** (DESY)
- *"PyTorch C++ API"*
by **David Brunner** (DESY)



Ivica's evening talk

Predicting the future




A Whale-Bus

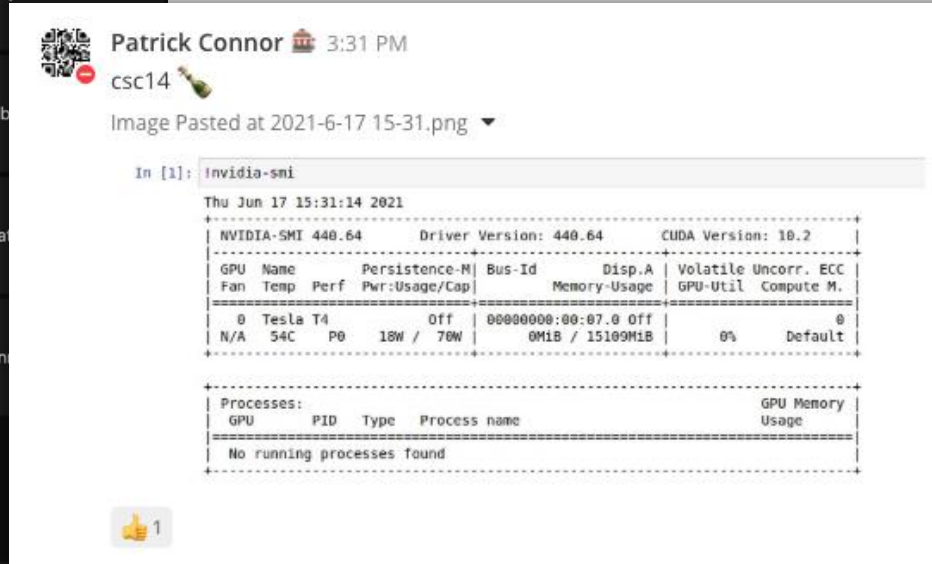
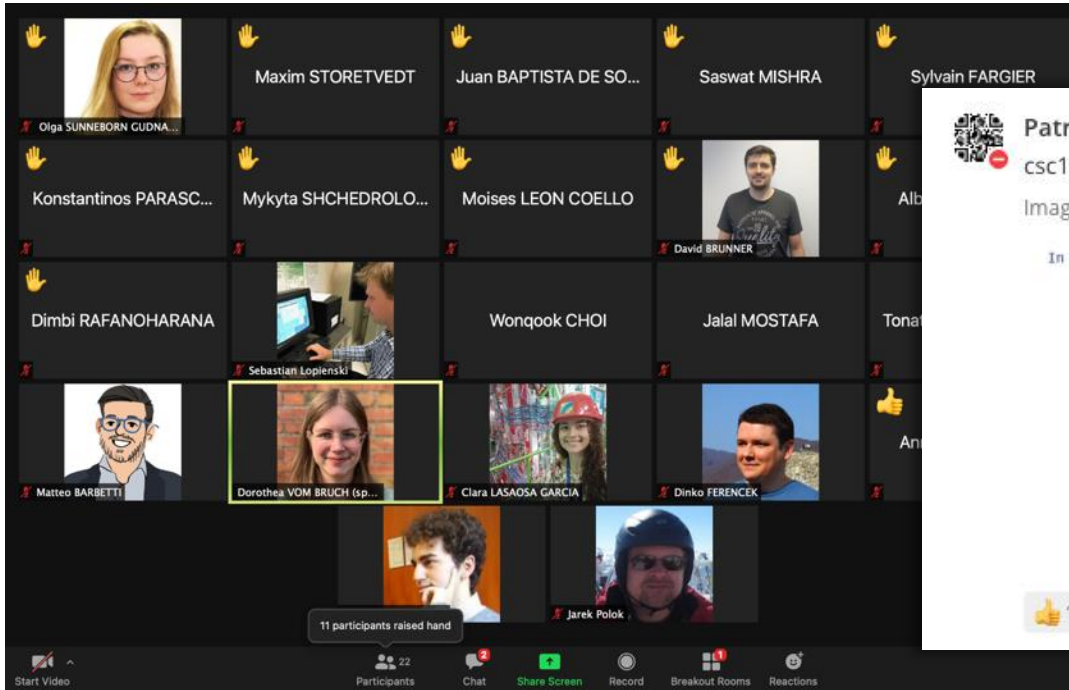


Hands-on exercises

(aka heating up CERN Data Centre 😊)



Hands-on exercises



Thanks for your patience during the Thursday's exercises;
and to CERN EP-SFT, IT-ST, IT-DB and IT-CM for providing GPUs via [SWAN](#) service

Technical issues? Last minute changes?

| Wednesday, 16 June 2021 | | Thursday, 17 June 2021 | |
|-------------------------|---|------------------------|--|
| 09:00 | Hardware evolution and heterogeneity - Andrzej Nowak | 09:00 | Data-oriented design - Andrzej Nowak |
| 10:00 | Practical vectorization - Sebastien Ponce (CERN) | 10:00 | Programming for GPUs - Dorothea vom Bruch (CPPM/CNRS) |
| 11:00 | Coffee break | 11:00 | Coffee break |
| 11:20 | Scientific computing on heterogeneous architectures - Dorothea vom Bruch (CPPM/CNRS) | 11:20 | Performant programming for GPUs - Daniel Campora (University of Maastricht) |



[Redacted] to Everyone 10:07 AM

Is it me or the sound is very bad?

[Redacted] to Everyone 10:08 AM

For me as well

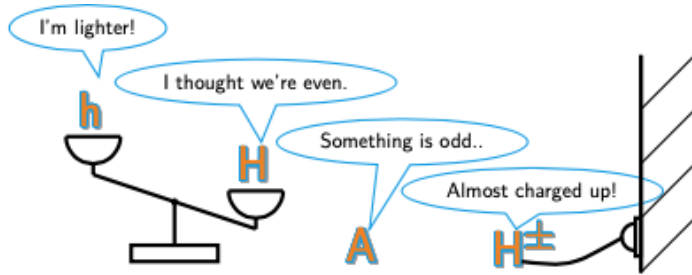
[Redacted] to Everyone 10:17 AM

The volume is too low for me. I can't get much :(

Self-presentation sessions:

one slide one minute per person

<https://indico.cern.ch/event/1017080/contributions/4291663/>



Read most of of Isaac Asimov's books

Not really proud that
that I'm always late

Cooking



Working here
(more or less)

Expectations



Reality



and I am a futsal referee.



"Remember to donate blood!"

Exam results

19 points out of 22

- Laura Martikainen *Helsinki Institute of Physics, Finland*
- Florian Bury *Catholic University of Louvain, Belgium*

20 points out of 22

- Marie Bachmayer *EPFL, Lausanne, Switzerland*
- Alberto Manfreda *INFN, Pisa, Italy*
- Konstantinos Paraschou *CERN; Aristotle University of Thessaloniki, Greece*
- Vincenzo Padulano *Valencia Polytechnic University, Spain*

22 points out of 22

- Olga Sunneborn Gudnadottir *Uppsala University, Sweden*

What now?

Anonymous feedback questionnaire

Please tell us what you think - help us get (even) better!

Survey

tCSC 2021 Spring Feedback

Lectures

For each lecturer, please evaluate their lectures with regard to the content, the quality of presentation, and the overall value.

Andrzej Nowak - Technologies and Platforms

Duration: 4 hours

Content *

Very poor

Poor

Fair

Good

Excellent

Three CERN Schools of Computing

iCSC 2020



Inverted school

16 Mar - 19 Mar 2020

[CERN](#) | Geneva | Switzerland

tCSC 2020



Thematic school

7 Jun - 13 Jun 2020

[MedILS](#) | Split | Croatia

CSC 2020



Main school

23 Aug - 5 Sep 2020

[University of Science and Technology \(AGH\)](#),
[and Institute of Nuclear Physics \(IFJ PAN\)](#) |
Kraków | Poland

Three CERN Schools of Computing



tCSC spring 2021

Thematic School 8th edition – online school –
14-18 June, 2021 – Registrations open: late-
March



tCSC autumn 2021

Thematic School 9th edition will take place in
Split, Croatia – 10-16 October 2021, further
information will be provided early June



CSC 2022

Main School 2022 43rd edition will take place
in Krakow (PL), 28 August – 10 September
2022, further information will be provided
early 2022

Attend other CSCs 😊

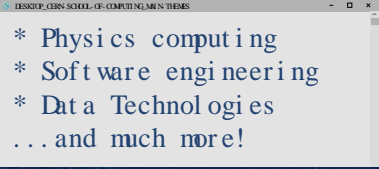
Advertise CSC
to your colleagues

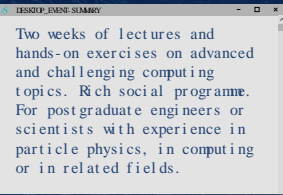


CERN School of Computing

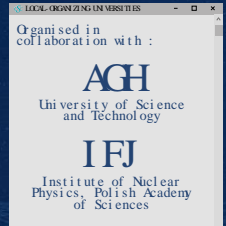
23rd August - 5th September 2021

Kraków - Poland

- 
- * Physics computing
 - * Software engineering
 - * Data Technologies
 - ... and much more!




Two weeks of lectures and hands-on exercises on advanced and challenging computing topics. Rich social programme. For postgraduate engineers or scientists with experience in particle physics, in computing or in related fields.



Organised in collaboration with:

AGH
University of Science and Technology

IFJ
Institute of Nuclear Physics, Polish Academy of Sciences





Back in
2005...

- ✓ Have you ever heard of **Enterprise Computing**, is it relevant to physics computing?
- ✓ Do you know what **Design Pattern** is?
- ✓ Are you sure the software you write has no **security holes**?
- ✓ Are you sure that you know and master **modern debugging tools**?
- ✓ Do you know how to design (effectively) a **database schema**?
- ✓ What is the secret to writing an efficient **SQL** query?
- ✓ What is database **performance tuning**, why is it perceived as magic and how to tame it?
- ✓ Do you know how to read an **execution plan**?
- ✓ How does **Google News** work?
- ✓ Do you know, in practice how to expose your application as a **Web Service**?
- ✓ Are you sure your **Web Services** are **secure**?

All the answers at ICSC



CERN
School of Computing

inverted CSC-2005
"Where students turn into teachers"



23-25 February 2005, CERN*

- ▶ Data Management and Data Bases
- ▶ Advanced Software Development and Engineering
- ▶ Web Services in Distributed Computing

- a novel idea prototyped in 2005
- a three-day series of lectures proposed and delivered by selected students
- advanced topics, rarely taught at CERN before

Lecturers - all former CSC2004 students

| | |
|----------------------------|--------------------------|
| Paolo Adragna | University of Siena |
| Miguel Anjo | CERN |
| Ioannis Baltopoulos | Imperial College |
| Gerhard Brandt | University of Heidelberg |
| Giovanni Chierico | CERN |
| Brice Copy | CERN, |
| Michal Kwiatek | CERN |
| Ruben Leivas Ledo | CERN |
| Sebastian Lopienski | CERN |
| Petr Olmer | CERN |
| Zornitsa Zaharieva | CERN |

<http://cern.ch/csc>

* IT Amphitheatre, building 31
Free attendance but registration recommended

"Where
students
turn into
teachers"

2019: the 12th edition of the iCSC

2005



2006



2008



2010



2011



2013




2014



2015



2016



2017




2018



2019



iCSC lecturers becoming full CSC lecturers



- Sebastian Lopienski (*iCSC 2005*)
- Andrzej Nowak (*iCSC 2008*)
- Benjamin Radburn-Smith (*iCSC 2010*)
- Thomas Keck (*iCSC 2016*)
- Eamonn Maguire (*iCSC 2017*)

Academic topics

(iCSC is not a conference!)



Artificial Intelligence, Machine Learning



Big Data, Hadoop, Apache Spark



Computational Physics, Tensor Networks



Container Orchestration



FPGAs (Introduction)



FPGAs + VHDL



Numerical Analysis, Finite Element Method



Pattern Recognition, Track Finding

iCSC 2019 schedule

<https://indico.cern.ch/event/766995/timetable>

| Monday, 4 March 2019 | Tuesday, 5 March 2019 | Wednesday, 6 March 2019 | Thursday, 7 March 2019 |
|--|--|---|---|
| 08:30 Welcome coffee | 08:30 Welcome coffee | 08:30 Welcome coffee | 08:30 Welcome coffee |
| 09:00 A word from the IT Department Head | 09:00 Hardware Acceleration Through FPGAs - Basics of VHDL (lecture 2) | 09:00 Tensor Networks - Singular Value Decomposition (exercise 1) | 09:00 Big Data Technologies and Physics Analysis with Apache Spark (exercise 1) |
| 09:10 Introduction to the Inverted CSC | | | |
| 09:15 A practical approach to Convolutional Neural Networks (lecture) | 10:00 Tensor Networks - Introduction and Matrix Product States (lecture 1) | 10:00 Tensor Networks - Application of the ITEBD Algorithm (exercise 2) | 10:00 Big Data Technologies and Physics Analysis with Apache Spark (exercise 2) |
| 10:15 Coffee | | | |
| 10:45 A Scientist's Guide to FPGAs | 11:00 Coffee | 11:00 Coffee | |
| 11:45 Lunch break / WIT Diversity Talk | 11:15 Tensor Networks - The ITEBD Algorithm (lecture 2) | 11:15 How container orchestration can strengthen your micro-services: the approach of Kubernetes (lecture) | |
| | 12:15 Lunch break | 12:15 Lunch break | |
| 13:30 Global track finding algorithms | 13:30 Big Data Technologies and Physics Analysis with Apache Spark (lecture 1) | 13:00 Data Centre visit | |
| 14:30 Hardware Acceleration Through FPGAs - Basic Concepts (lecture 1) | 14:30 Big Data Technologies and Physics Analysis with Apache Spark (lecture 2) | 14:30 Efficient C++ implementation of custom FEM kernel with Eigen | |
| 15:30 Coffee | 15:30 Coffee | 15:30 Coffee | |
| 16:00 A practical approach to Convolutional Neural Networks (exercise 1) | 16:00 Hardware Acceleration Through FPGAs - First Experiments in VHDL (exercise 1) | 16:00 How container orchestration can strengthen your micro-services: the approach of Kubernetes (exercise 1) | |
| 17:00 A practical approach to Convolutional Neural Networks (exercise 2) | 17:00 Hardware Acceleration Through FPGAs - Easy DSP Applications (exercise 2) | 17:00 How container orchestration can strengthen your micro-services: the approach of Kubernetes (exercise 2) | |

Lectures in
IT
Amphitheat
Exercises

in 513-1-
024

Take your
laptop

iCSC 2019: a **record** attendance

- 400+ registrations from ~90 institutes in 30 countries
 - 164 with ...@cern.ch e-mail addresses; 29 with others
 - 20% females
 - IT department (44 people), BE department (36), CMS (24), ATLAS (14)
- 60 people at the opening (+ ~25 by webcast)
 - 30-60 people for each lecture
- Some people travelled to Geneva (by plane etc.) to attend iCSC!

"This is how you fish"

The slide content includes:

The barrel, the fish, the rod...

```
public class ShapeService {
    private ApplicationContext pool;

    public void init() {
        // create pool
        pool = Spring.run
            (ApplicationConfig.class);
    }

    public double findArea(String shape, ...) {
        // return the shape
        Shape fishedOut = pool.getBean(shape);
        return fishedOut.findArea(...);
    }
}

public interface Shape {
    public double findArea(...);
}

public class ApplicationConfig {
    @Bean
    public Shape rectangle() {
        return new Rectangle(...);
    }

    @Bean
    @Lazy
    public Shape circle() {
        return new Circle(...);
    }
}

public class Rectangle implements Shape {
    public double findArea(...){
        return ...;
    }
}
```

Diagram: A fish labeled 'Seabass' is being reeled in by a rod and reel. The fish is in a blue pool. An arrow points from the fish to the text 'Fish'.





Bitcoin coffee discussions








Inverted CSC (iCSC) 2022?

- **October 2021** - students of tCSC spring 2021 and tCSC autumn 2021 sent their lecture and exercise proposals
 - what will you present? 😊
- **November 2021** – go/no-go for the iCSC 2022
- **mid March 2022** – iCSC 2022

Thanks

Danilo, Andrzej, Sebastien, Arthur, Dorothea, Daniel and Ivica



Joelma



Jarek



Thank YOU, tCSC 2021 participants



