



DEEP
LEARNING
INSTITUTE



PRACE Workshop: Deep Learning and GPU Programming Workshop

10 – 13 May 2022



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Overview



- The workshop is co-organised by LRZ, CSC, IT4Innovations and NVIDIA Deep Learning Institute (DLI) for the Partnership for Advanced Computing in Europe (PRACE).
- NVIDIA Deep Learning Institute (DLI) offers hands-on training for developers, data scientists, and researchers looking to solve challenging problems with deep learning.
- The online workshop combines lectures about Accelerated Computing with OpenACC and CUDA with lectures about Fundamentals of Deep Learning for single and for Multi-GPUs.
- Learn how to accelerate your applications with OpenACC and CUDA, how to train and deploy a neural network to solve real-world problems, and how to effectively parallelize training of deep neural networks on Multi-GPUs.
- The lectures are interleaved with many hands-on sessions using Jupyter Notebooks. The exercises will be done on a fully configured GPU-accelerated workstation in the cloud.

PRACE Training



LRZ as part of the Gauss Centre for Supercomputing (GCS), CSC and IT4Innovations belong to the 14 **PRACE Training Centres** that started in 2012-2017-2020:

- Barcelona Supercomputing Center (Spain)
- CINECA Consorzio Interuniversitario (Italy)
- CSC – IT Center for Science Ltd (Finland)
- EPCC at the University of Edinburgh (UK)
- Gauss Centre for Supercomputing (Germany)
- Maison de la Simulation (France)
- GRNET – Greek Research and Technology Network (Greece)
- ICHEC – Irish Centre for High-End Computing (Ireland)
- IT4I – National Supercomputing Center VSB Technical University of Ostrava (Czech Republic)
- SURFsara (The Netherlands)
- TU Wien – VSC Research Center (Austria)
- University ANTWERPEN – VSC & CÉCI (Belgium)
- University of Ljubljana – HPC Center Slovenia (Slovenia)
- Swedish National Infrastructure for Computing (SNIC) (Sweden)



Univerza v Ljubljani



Irish Centre for High-End Computing

VSB TECHNICAL UNIVERSITY OF OSTRAVA

IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER



Mission: Serve as **European hubs and key drivers of advanced high-quality training** for researchers working in the computational sciences.

<http://www.training.prace-ri.eu/>





DEEP
LEARNING
INSTITUTE



DEEP LEARNING INSTITUTE

DLI Mission: Help the world to solve the most challenging problems using AI and deep learning

We help developers, data scientists and engineers to get started in architecting, optimizing, and deploying neural networks to solve real-world problems in diverse industries such as autonomous vehicles, healthcare, robotics, media & entertainment and game development.



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Lecturers



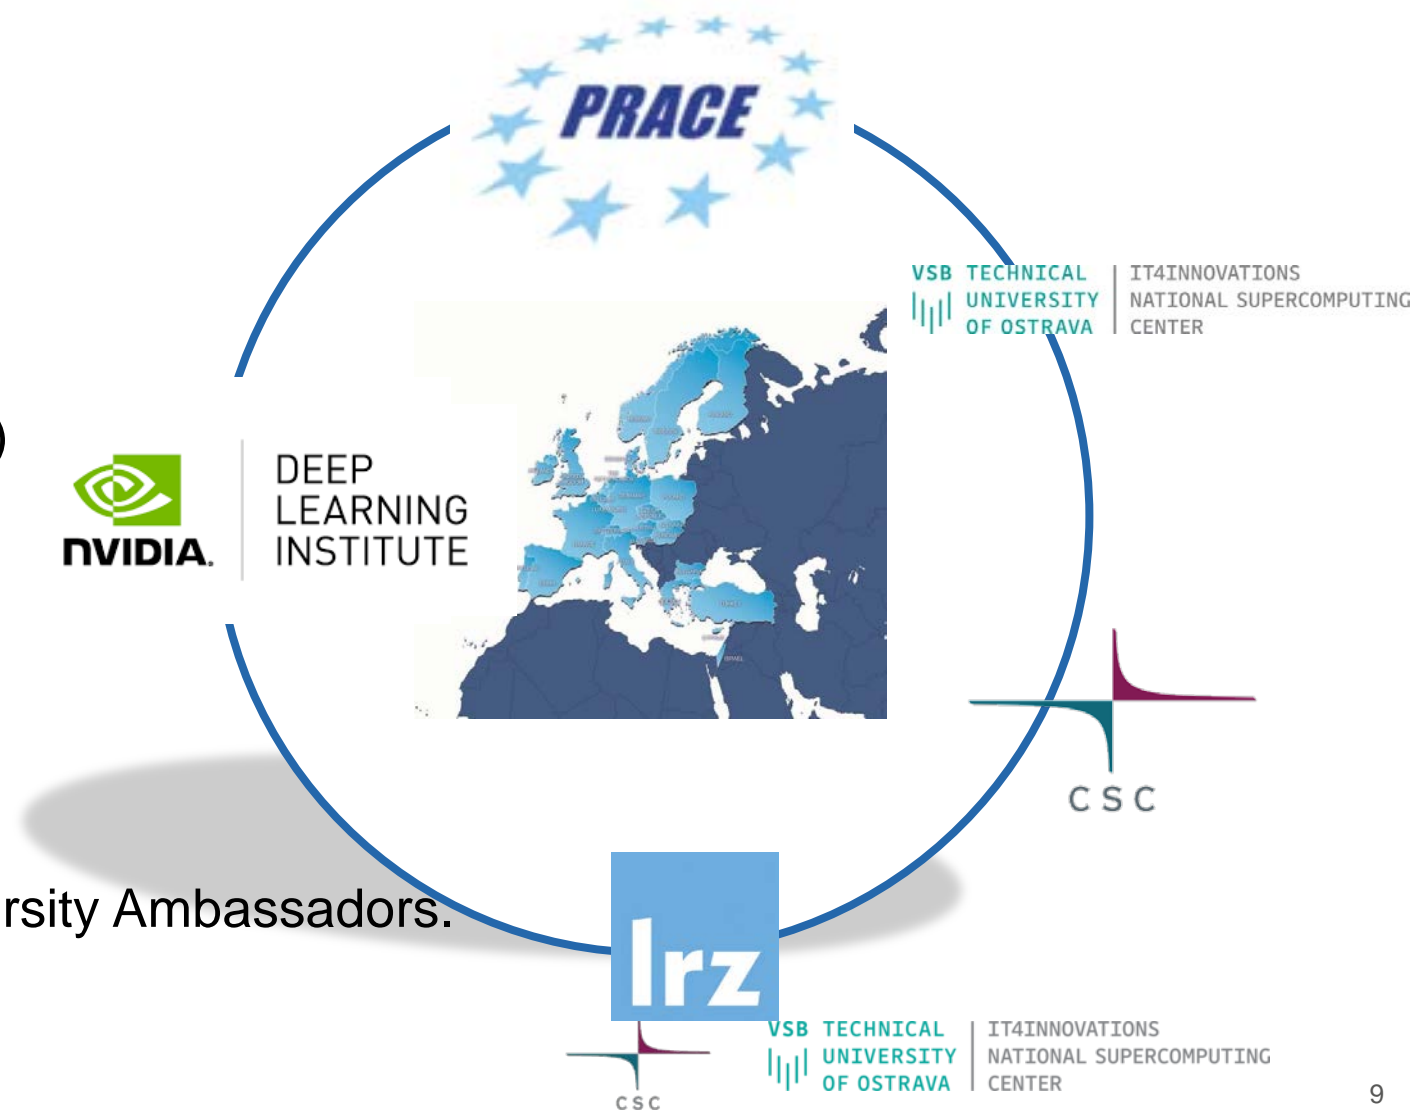
- **Lecturers:**

- Dr. Momme Allalen (LRZ)
- PD Dr. Juan Durillo Barrionuevo (LRZ)
- Dr. Volker Weinberg (LRZ)
- Georg Zitzlsberger (IT4Innovations)

All instructors are NVIDIA certified University Ambassadors.



DEEP
LEARNING
INSTITUTE



1st day: Fundamentals of Accelerated Computing with OpenACC



- On the 1st day you learn the basics of OpenACC, a high-level programming language for programming on GPUs. Discover how to accelerate the performance of your applications beyond the limits of CPU-only programming with simple pragmas.
- You'll learn:
 - How to profile and optimize your CPU-only applications to identify hot spots for acceleration
 - How to use OpenACC directives to GPU accelerate your codebase
 - How to optimize data movement between the CPU and GPU accelerator
- Upon completion, you'll be ready to use OpenACC to GPU accelerate CPU-only applications.

2nd day: Fundamentals of Accelerated Computing with CUDA C/C++



- The CUDA computing platform enables the acceleration of CPU-only applications to run on the world's fastest massively parallel GPUs. On the 2nd day you experience C/C++ application acceleration by:
 - Accelerating CPU-only applications to run their latent parallelism on GPUs
 - Utilizing essential CUDA memory management techniques to optimize accelerated applications
 - Exposing accelerated application potential for concurrency and exploiting it with CUDA streams
 - Leveraging command line and visual profiling to guide and check your work
- Upon completion, you'll be able to accelerate and optimize existing C/C++ CPU-only applications using the most essential CUDA tools and techniques. You'll understand an iterative style of CUDA development that will allow you to ship accelerated applications fast.

3rd day: Fundamentals of Deep Learning



- Explore the fundamentals of deep learning by training neural networks and using results to improve performance and capabilities.
- During this day, you'll learn the basics of deep learning by training and deploying neural networks. You'll learn how to:
 - Implement common deep learning workflows, such as image classification and object detection
 - Experiment with data, training parameters, network structure, and other strategies to increase performance and capability
 - Deploy your neural networks to start solving real-world problems
- Upon completion, you'll be able to start solving problems on your own with deep learning.

4th day: Fundamentals of Deep Learning for Multi-GPUs



- The computational requirements of deep neural networks used to enable AI applications like self-driving cars are enormous. A single training cycle can take weeks on a single GPU or even years for larger datasets like those used in self-driving car research. Using multiple GPUs for deep learning can significantly shorten the time required to train lots of data, making solving complex problems with deep learning feasible.
- On the 4th day we will teach you how to use multiple GPUs to train neural networks. You'll learn:
 - Approaches to multi-GPUs training
 - Algorithmic and engineering challenges to large-scale training
 - Key techniques used to overcome the challenges mentioned above
- Upon completion, you'll be able to effectively parallelize training of deep neural networks using TensorFlow.

Tentative Agenda Day 1: **Fundamentals of Accelerated Computing with OpenACC**



- 10:00-10:15 Intro
- 10:15-12:00 Profiling
- **12:00-13:00 Lunch Break**
- 13:00-14:20 OpenACC Directives
- **14:20-14:30 Coffee Break**
- 14:30-15:45 GPU Programming and Data Management
- 15:45-16:00 Q&A, Final Remarks




All times are in EEST

Tentative Agenda Day 2: Fundamentals of Accelerated Computing with CUDA C/C++



- 10:00-10:15 Introduction CUDA C/C++
- 10:15-12:00 Accelerating Applications with CUDA C/C++

- **12:00-13:00 Lunch Break**



All times are in EEST

- 13:00-14:20 Managing Accelerated Application Memory with CUDA unified memory and nsys
- **14:20-14:30 Coffee Break**
- 14:30-15:45 Asynchronous Streaming and Visual Profiling for Accelerated Applications with CUDA C/C++
- 15:45-16:00 Q&A, Final Remarks


Tentative Agenda Day 3: Fundamentals of Deep Learning



10:00-10:20 Welcome and Intro

10:20-12:00 Introduction to Deep Learning and Convolutional Neural Networks

12:00-13:00 Lunch Break



All times are in EEST

13:00-14:20 Data Augmentation, Deployment and Pre-Trained Models

14:20-14:30 Coffee Break

14:30-15:45 Advanced Architectures

15:45-16:00 Q&A

Tentative Agenda Day 4: **Fundamentals of Deep Learning for Multi-GPUs**



10:00-10:15 Introduction
10:15-12:00 Stochastic Gradient Descent


12:00-13:00 Lunch Break

13:00-14:20 Introduction to Distributed Training

14:20-14:30 Coffee Break

14:30-15:45 Algorithmic Challenges of Distributed SGD

15:45-16:00 Q&A



All times are in EEST

Workshop Webpage



DEEP
LEARNING
INSTITUTE



- All slides will be made available during the workshop under:

- <https://tinyurl.com/dli-workshop-csc>

- Further information on:

- Agenda
- Training Setup
- Slides
- Documentation



Training Setup



- To get started, follow these steps:
- Create an NVIDIA Developer account at <http://courses.nvidia.com/join> Select "Log in with my NVIDIA Account" and then "Create Account".
- If you use your own laptop, make sure that WebSockets works for you:
Test your Laptop at <http://websocketstest.com>
 - Under ENVIRONMENT, confirm that "WebSockets" is checked yes.
 - Under WEBSOCKETS (PORT 80]. confirm that "Data Receive", "Send", and "Echo Test" are checked yes.
 - If there are issues with WebSockets, try updating your browser.
We recommend Chrome, Firefox, or Safari for an optimal performance.
- Visit <http://courses.nvidia.com/dli-event> and enter the event code provided by the instructor.
- You're ready to get started.

- To ensure a pleasant experience with Zoom Meeting, we encourage participants to **download and install the latest Zoom application** via <https://zoom.us/download>.
- If you have problems with your computer audio, you can also **join by phone**. Find your local number: <https://us02web.zoom.us/j/kdKozJPtKk>

Zoom Meeting

00:12:30 View

Switch between Speaker View / Gallery View / Fullscreen Mode

Speaker View

✓ Gallery View

Fullscreen

Start / Stop Video / Select Virtual Background

Un-/Mute

Unmute Start Video Participants Chat Share Screen Record Reactions Leave

The screenshot shows the Zoom meeting interface. A red arrow points from the text 'Switch between Speaker View / Gallery View / Fullscreen Mode' to the 'View' dropdown menu, which is open and shows 'Speaker View', 'Gallery View' (selected), and 'Fullscreen'. Another red arrow points from the text 'Un-/Mute' to the 'Unmute' button in the bottom toolbar. A third red arrow points from the text 'Start / Stop Video / Select Virtual Background' to the 'Start Video' button in the bottom toolbar. The bottom toolbar also includes 'Participants', 'Chat', 'Share Screen', 'Record', 'Reactions', and 'Leave' buttons.

zoom

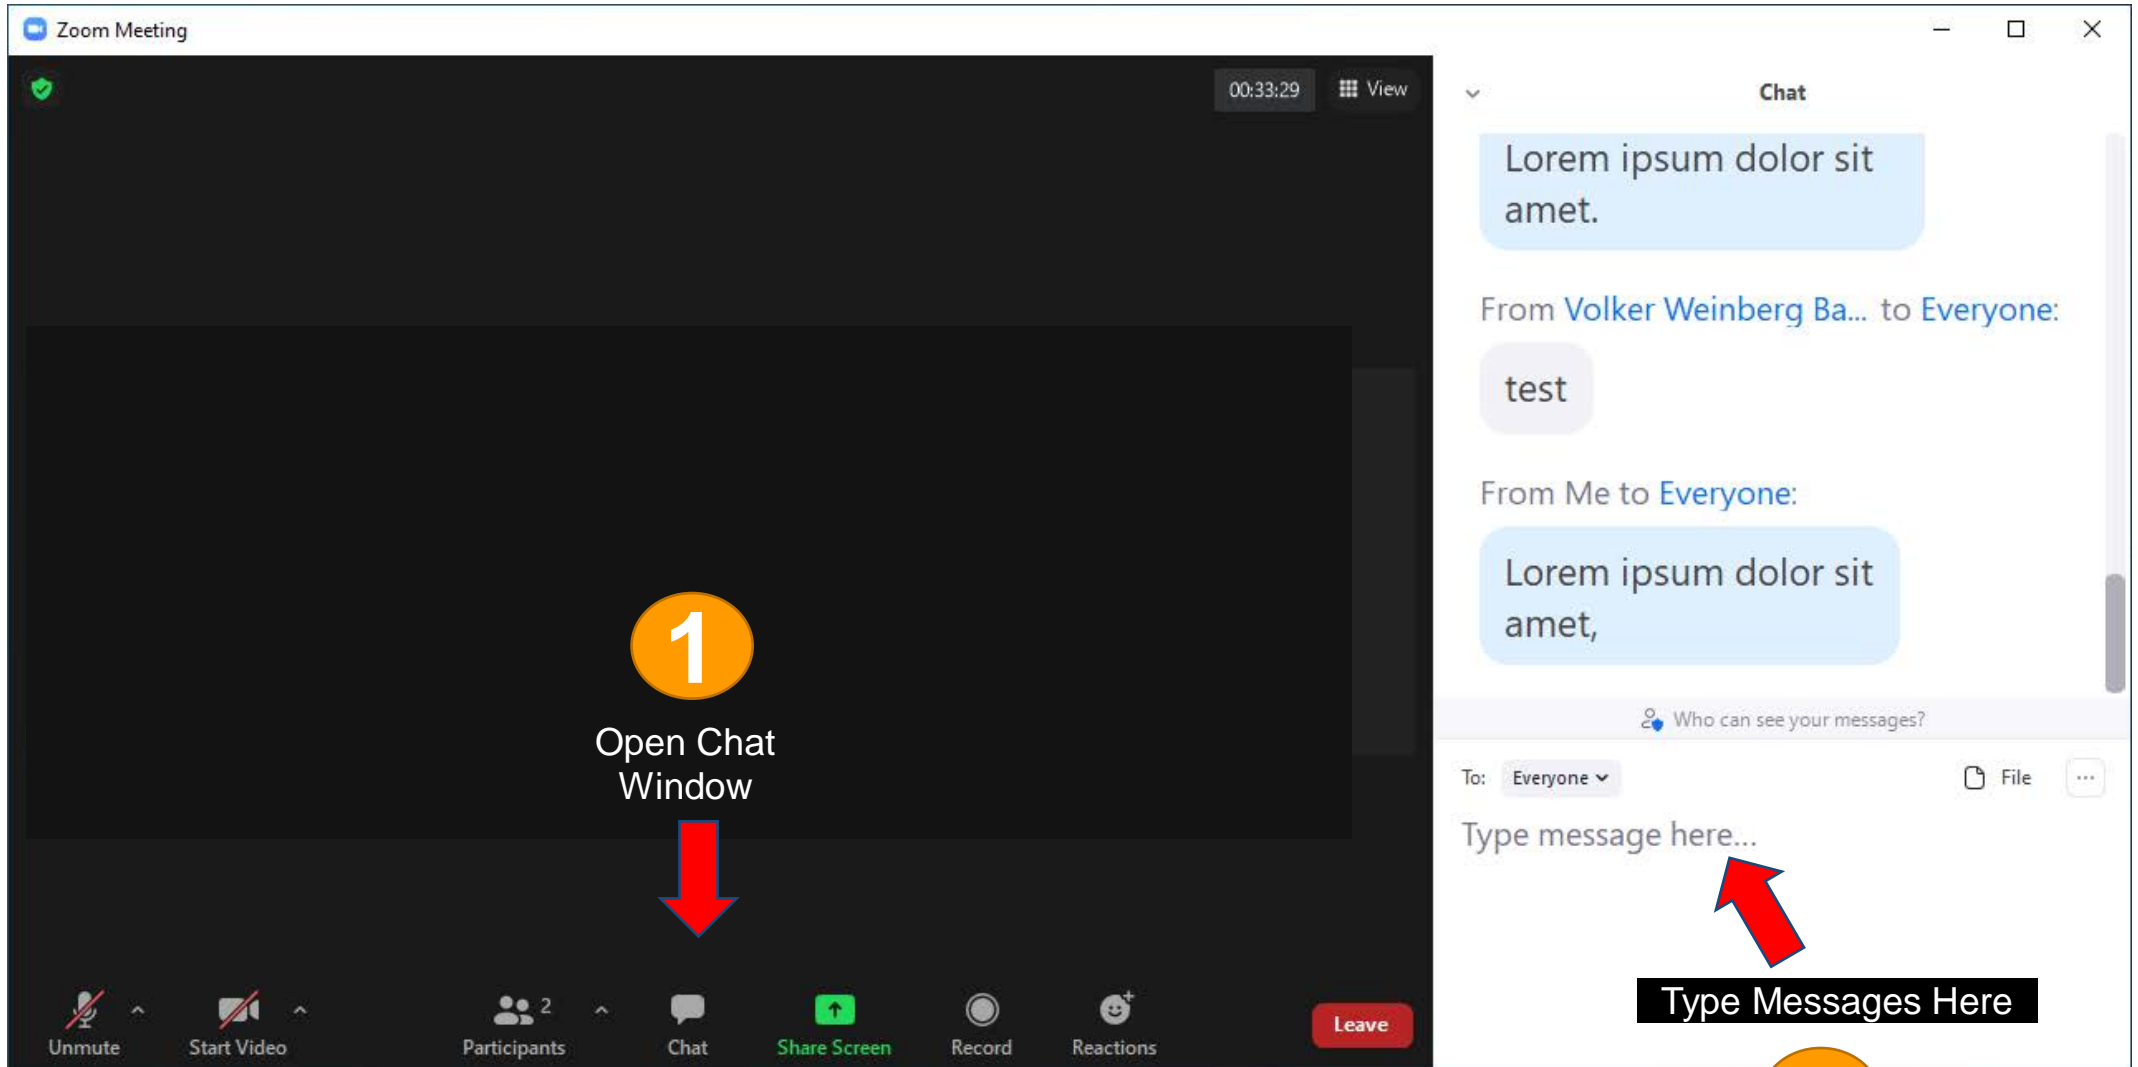
Participants List



Kindly use “<first name> <last name> (<institute>)” as your screenname.

Otherwise you will not receive a certificate of attendance after the course.

The screenshot shows a Zoom Meeting window with the Participants List open. The list contains two participants: Volker Wei... (Me) and Volker Weinberg Backup (Host). A context menu is open over the 'Me' participant, showing options: Edit Profile Picture, Rename, Unmute, and More >. A red arrow points from the 'Rename Yourself' text to the 'Rename' option. A large orange circle with the number '1' is overlaid on the 'Participants' button in the bottom toolbar, with the text 'Open Participants List' and a red arrow pointing down to the button. A second large orange circle with the number '2' is overlaid on the 'Rename' option in the context menu, with the text 'Rename Yourself' and a red arrow pointing to it.



Zoom Meeting

00:33:29 View

Chat

Lorem ipsum dolor sit amet.

From Volker Weinberg Ba... to Everyone:

test

From Me to Everyone:

Lorem ipsum dolor sit amet,

Who can see your messages?

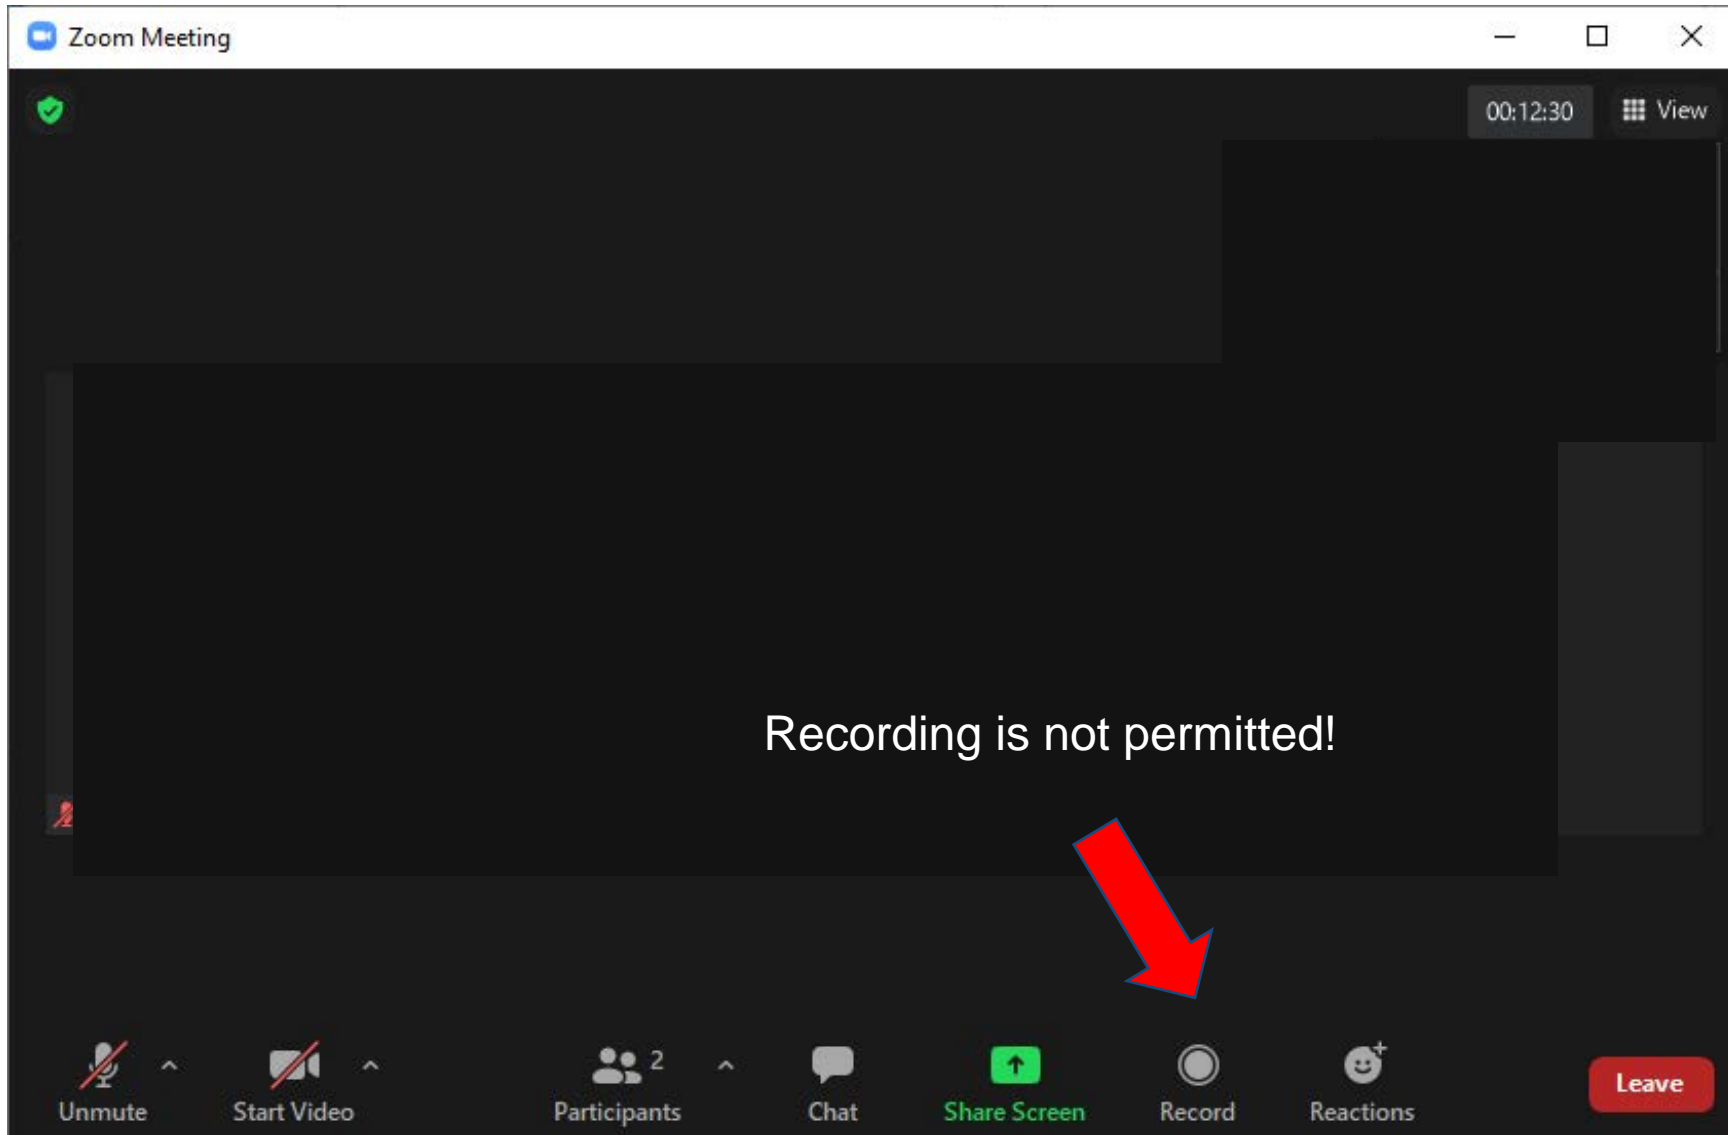
To: Everyone

Type message here...

Unmute Start Video Participants Chat Share Screen Record Reactions Leave


1
Open Chat Window

2
Type Messages Here



- Use **chat window** to ask questions **during the presentations**.
- In **Q&A sessions**:
 - You can also **raise your hand** if you have questions.
 - If you do not mind, please **show your video when asking questions** to make this workshop as interactive as possible.
- **Push to Talk**: The Push to Talk feature allows you to remain muted throughout the Zoom meeting and only if you hold down the spacebar you will be unmuted.

- **Instant Feedback:**



And now ...



DEEP
LEARNING
INSTITUTE



Enjoy the workshop!



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER