Good Afternoon!

Good Evening!

Good Night!

Good Late Night!

Good early Morning!

Good Morning!

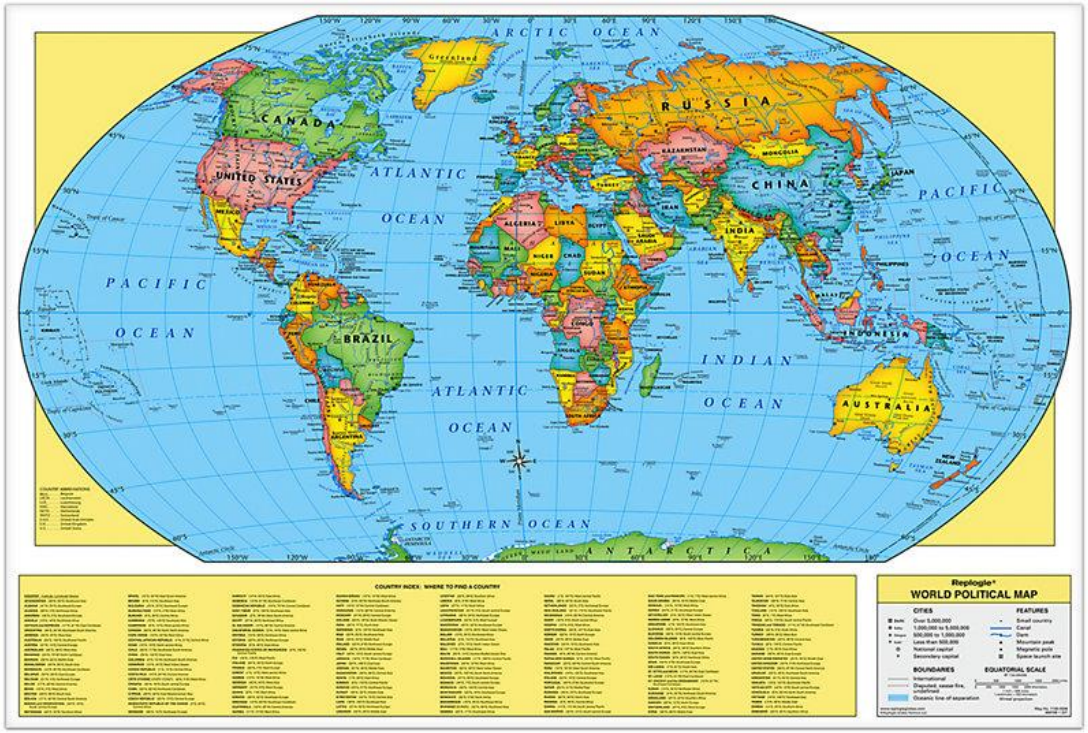

Thank you for joining the 6th

European 

workshop

(year 36 of the Condor project)

(year 14 of the CHTC)

State of Distributed High Throughput Computing

Miron Livny

John P. Morgridge Professor of Computer Science

Director of UW-Madison Center for High Throughput Computing

Technical director of the Open Science Grid

I am here today to report on the state of distributed High Throughput Computing - **dHTC**.

It is my privilege to report that **dHTC** is stronger than ever and is positioned to continue and get stronger in the short and long term. It builds on recent recognition (including a National Academy of Science (NAS) report!) of the value of **HTC** (some refer to it as capacity computing) to research and the growing adoption of distributed computing.

Moving forward!

Two weeks ago we were notified by the **NSF** that our proposal titled "**Partnership for Advanced Throughput Computing (PATh)**" was fully funded for five years with an annual budget of $4.5M

> *"The Partnership to Advance Throughput Computing (PATh) project will expand Distributed High Throughput Computing (dHTC) technologies and methodologies through innovation, translational effort, and large-scale adoption to advance the Science & Engineering goals of the broader community."*

www.nsf.gov/awardsearch/showAward?AWD_ID=2030508

PATh - a powerful partnership

- PATh is a partnership between the UW-Madison Center for High Throughput Computing (**CHTC**) and the Open Science Grid (**OSG**) Consortium
- PATh has two main elements: the **HTCondor Software Suite (HTCSS)** and the **Fabric of Capacity Services (FoCaS)**
- PATh ledership: Brian Bockelman (Co-PI), Miron Livny (PI), Lauren Michael (Co-PI), Todd Tannenbaum (Co-PI),Frank Wuerthwein (Co-PI)

Open Science Grid is 15 Years Old!

The OSG Story

Many factors contributed to the formation of the OSG consortium in 2005 and the initial DoE and NSF funding in 2006.

- The word Grid was hot
- Grid-3 was a success,
- US-LHC needed a distributed infrastructure.

The 2006 proposal presented the OSG as **a PetaScale Nationwide Facility** with LHC and LIGO as the primary physics drivers. The purpose was to offer US-LHC and LIGO what they want/need.

OSG and the HTC vision

For some of us OSG was a unique opportunity in a long quest to bring the power of distributed computing and the methodologies and technologies of high throughput computing (**HTC**) to the open science community. Far beyond large physics experiments and Grids.

In 1992 we presented at **CERN** our vision for a worldwide **flock of Condor pools**, in 1996 we introduced the concept of **HTC**, in 1997 we formulated the vision for a **National Technology Grid** for the NCSA Alliance, and in 2003 we deployed the Grid Laboratory of Wisconsin (**GLOW**) on the UW-Madison campus.

Dear Matt, Michael and Miron,

To celebrate the 40 years anniversary of the International Conference on Distributed Systems (ICDCS), the IEEE Technical Committee on Distributed Processing (TCDP) Award Committee has selected ICDCS High Impact Papers from the conference papers in the last 40 years, which have profoundly influenced the field of distributed systems. The high impact paper awards are sponsored jointly by ICDCS 2020 and IEEE TCDP. The Awards will be presented in ICDCS 2020 in Singapore in December 2020.

The committee decided to award the following two ICDCS High Impact Papers to

John Ousterhout,
for his paper entitled "Scheduling Techniques for Concurrent Systems" in ICDCS 1982;

and to

Michael J. Litzkow, Miron Livny and Matt W. Mutka
for their paper entitled "Condor – A Hunter of Idle Workstations" in ICDCS 1988.

The author(s) of each paper will receive a cash award of $2,000. We will contact you again for other arrangements related to the award.

Congratulations!

IEEE TCDP Award Committee

Ken Birman, Cornell University
Jiannong Cao (co-Chair), Hong Kong Polytechnic University
Chita Das, Pennsylvania State University
Nancy Lynch, Massachusetts Institute of Technology
Michel Raynal, IRISA, University of Rennes
Jie Wu, Temple University
Zheng-Zhong Xu, Macau University
Xiaodong Zhang (co-Chair), Ohio State University

Condor - A Hunter of Idle Workstations

Michael J. Litzkow, Miron Livny, and Matt W. Mutka

Department of Computer Sciences
University of Wisconsin
Madison, WI 53706

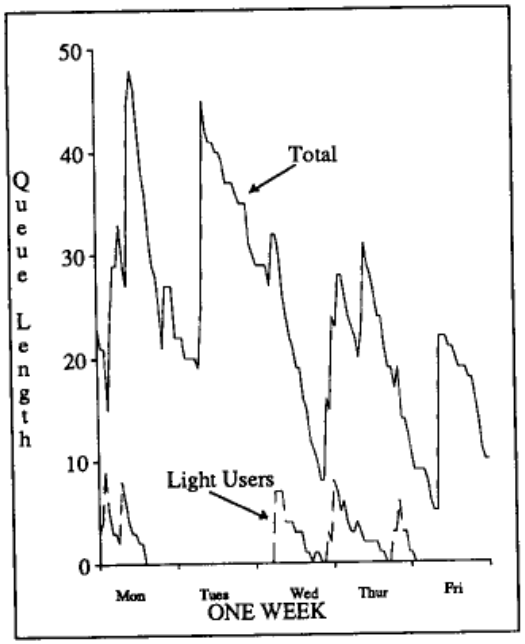

Figure 7: Queue Lengths for One Week.

There are over 100 VAXstation II workstations in our department. Since less than 30% of their capacity is utilized [1], a system has been designed and implemented to execute jobs remotely at idle workstations. Within our department there are many users working on problems that need large amounts of computing capacity. A few example problems include studies of load-balancing algorithms [9], simulation of real-time scheduling algorithms [10], studies of neural network learning models [11], and mathematical combinatorial problems [12]. These jobs typically require several hours of CPU time and little interaction with their users. The Condor system is designed to serve these users by executing their long running background jobs at idle workstations. To make our system attractive to these users, several issues must be addressed. First, the placement of background jobs should be transparent to users. The system should be responsible for knowing when workstations are idle and users should not need to know where their remote jobs execute. Second, if a remote site running a background job fails, the job should be restarted automatically at some other location to guarantee job completion. Third, since a workstation can serve as a source of remote cycles for others when it is not used by its owner, users expect to receive fair access to cycles when remote capacity is wanted. Fourth, the mechanisms implementing the system are expected to consume very little capacity. Otherwise users would not allow their workstations to be part of such a system if it interferes with their local activity.

OSG adopts dHTC

The **dHTC** acronym was first introduced about a decade ago in the 2011 proposal for the second OSG phase that was titled "**THE OPEN SCIENCE GRID The Next Five Years: Distributed High Throughput Computing for the Nation's Scientists, Researchers, Educators, and Students".**

- The abstract stated that ***"The Open Science Grid is a distributed computing infrastructure for large-scale scientific research. The OSG contributes to the Worldwide LHC Computing Grid as the shared distributed computing facility used by the US ATLAS and US CMS experiments***."
- The body of the proposal under project vision stated – "***The members of the OSG consortium are united in a commitment to promote the adoption and to advance the state of the art of distributed high throughput computing (DHTC) – shared utilization of autonomous resources where all the elements are optimized for maximizing computational throughput."***

OSG “Statement of Purpose”

OSG is a consortium dedicated to the advancement of all of open science via the practice of distributed High Throughput Computing (dHTC), and the advancement of its state of the art.

Distributed is “in”!

Why do we need the “**d**” since HTC was from the very beginning based on the principals of distributed systems?

- Part of the reason is that “distributed” is in these days. The 2020 **NSF-Campus Cyberinfrastructure (CC*)** solicitation used “distributed” more than 15 times.
- Distributed technologies like DASK and K8S are widely adopted fueled by the distributed nature of commercial clouds.

CHTC and OSG are natural partners

- The CHTC and OSG approach to **dHTC** is strongly aligned with the vision driving the NSF-CC* program as it follows a strong **campus focus** – researchers are on **campuses** resources are on **campuses** and they both recognize the value of sharing across campus boundaries.
- CHTC has pioneered research computing facilitation, provides HTC services across the UW campus, is an active member of the OSG and is home of the HTCSS. This suite plays an import (and growing) role in the fabric of dHTC services that OSG provides.

The transition from the Condor System to the HTCondor Software Suite (HTCSS)

- Software elements of the **HTCSS** can be used "stand alone" or "mix and match"
- On-the-fly and hot deployment and upgrades of **HTCSS** elements
- **HTCSS** interfaces with "other/external" technologies, execution environments and services
- **HTCSS** leverages functionality of widely adopted tools

Running a 51k GPU burst for Multi-Messenger Astrophysics with IceCube across all available GPUs in the Cloud

Frank Würthwein - OSG Executive Director
Igor Sfiligoi - Lead Scientific Researcher
UCSD/SDSC

Jensen Huang keynote @ SC19

The Largest Cloud Simulation in History

Saturday morning before SC19 we bought all GPU capacity that was for sale in Amazon Web Services, Microsoft Azure, and Google Cloud Platform worldwide

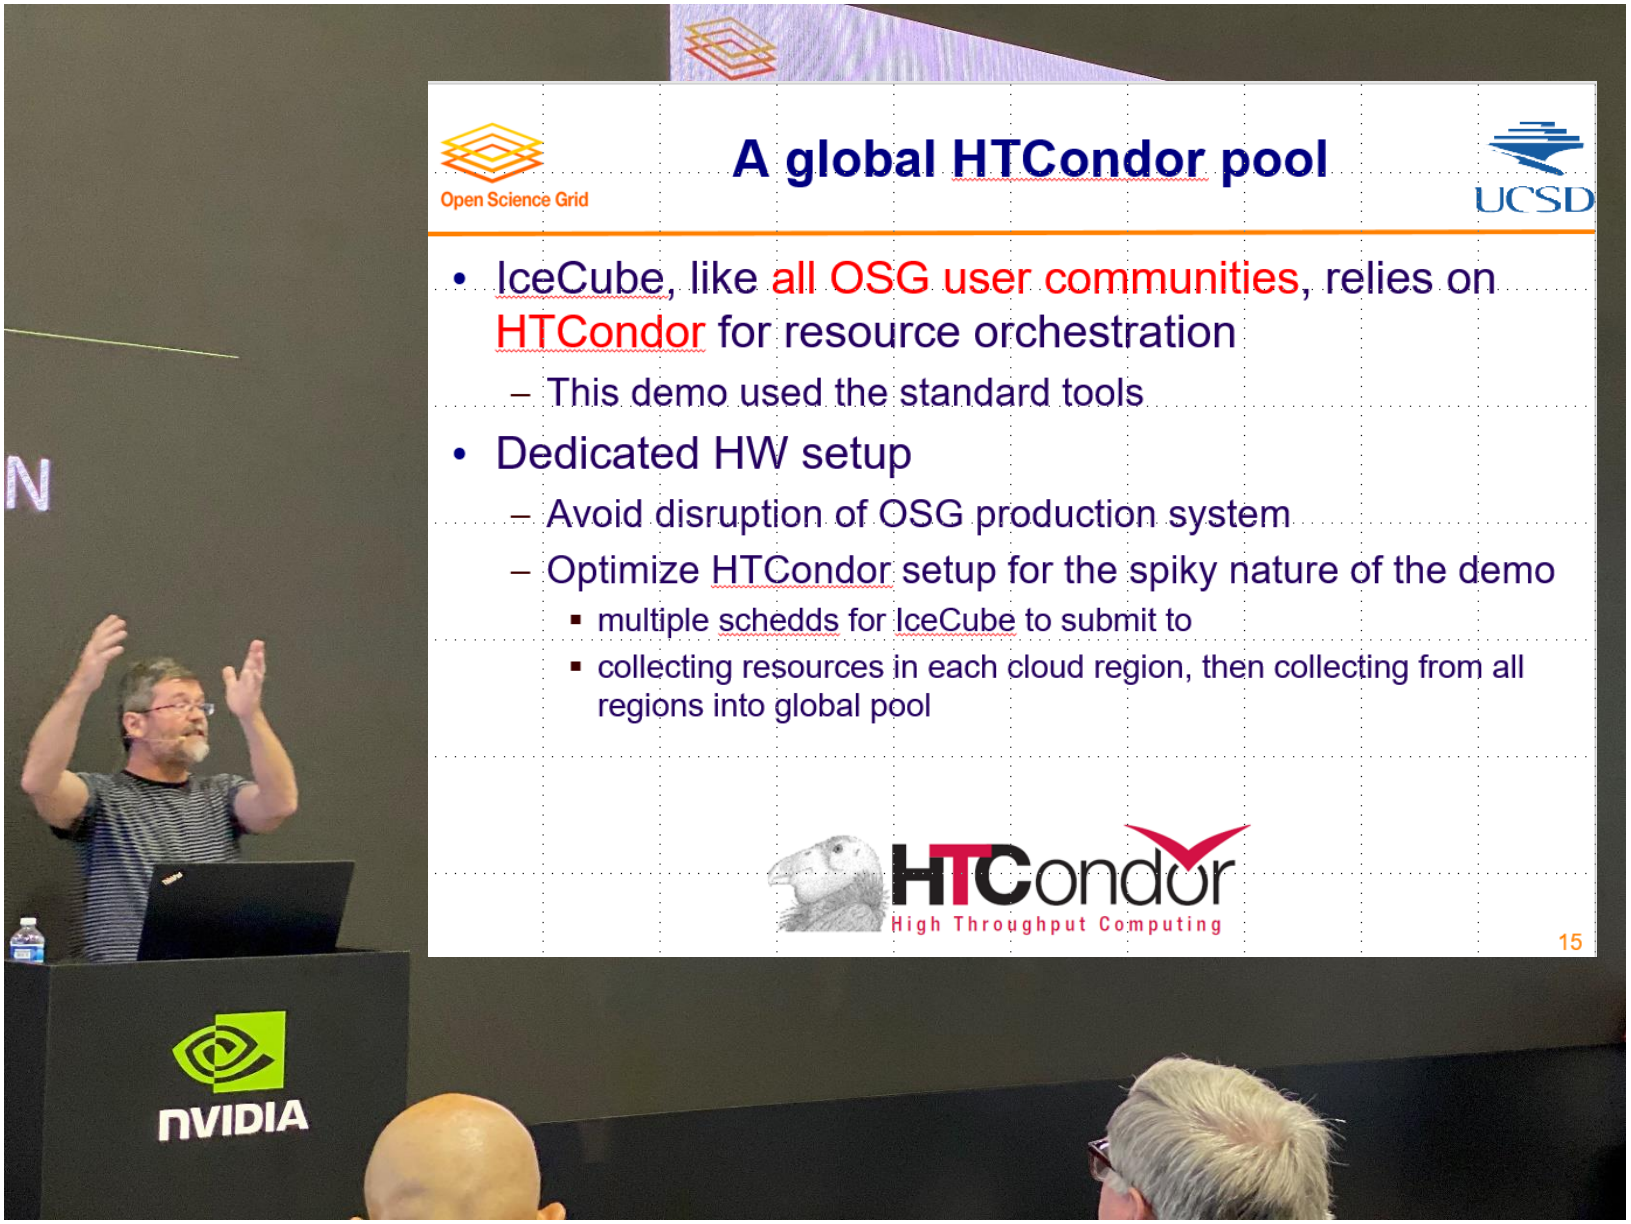
Open Science Grid
A global HTCondor pool
UCSD
• IceCube, like all OSG user communities, relies on HTCondor for resource orchestration
– This demo used the standard tools
• Dedicated HW setup
– Avoid disruption of OSG production system
– Optimize HTCondor setup for the spiky nature of the demo
▪ multiple schedds for IceCube to submit to
▪ collecting resources in each cloud region, then collecting from all regions into global pool
HTCondor
High Throughput Computing
15
NVIDIA

HTCondor Distributed CI

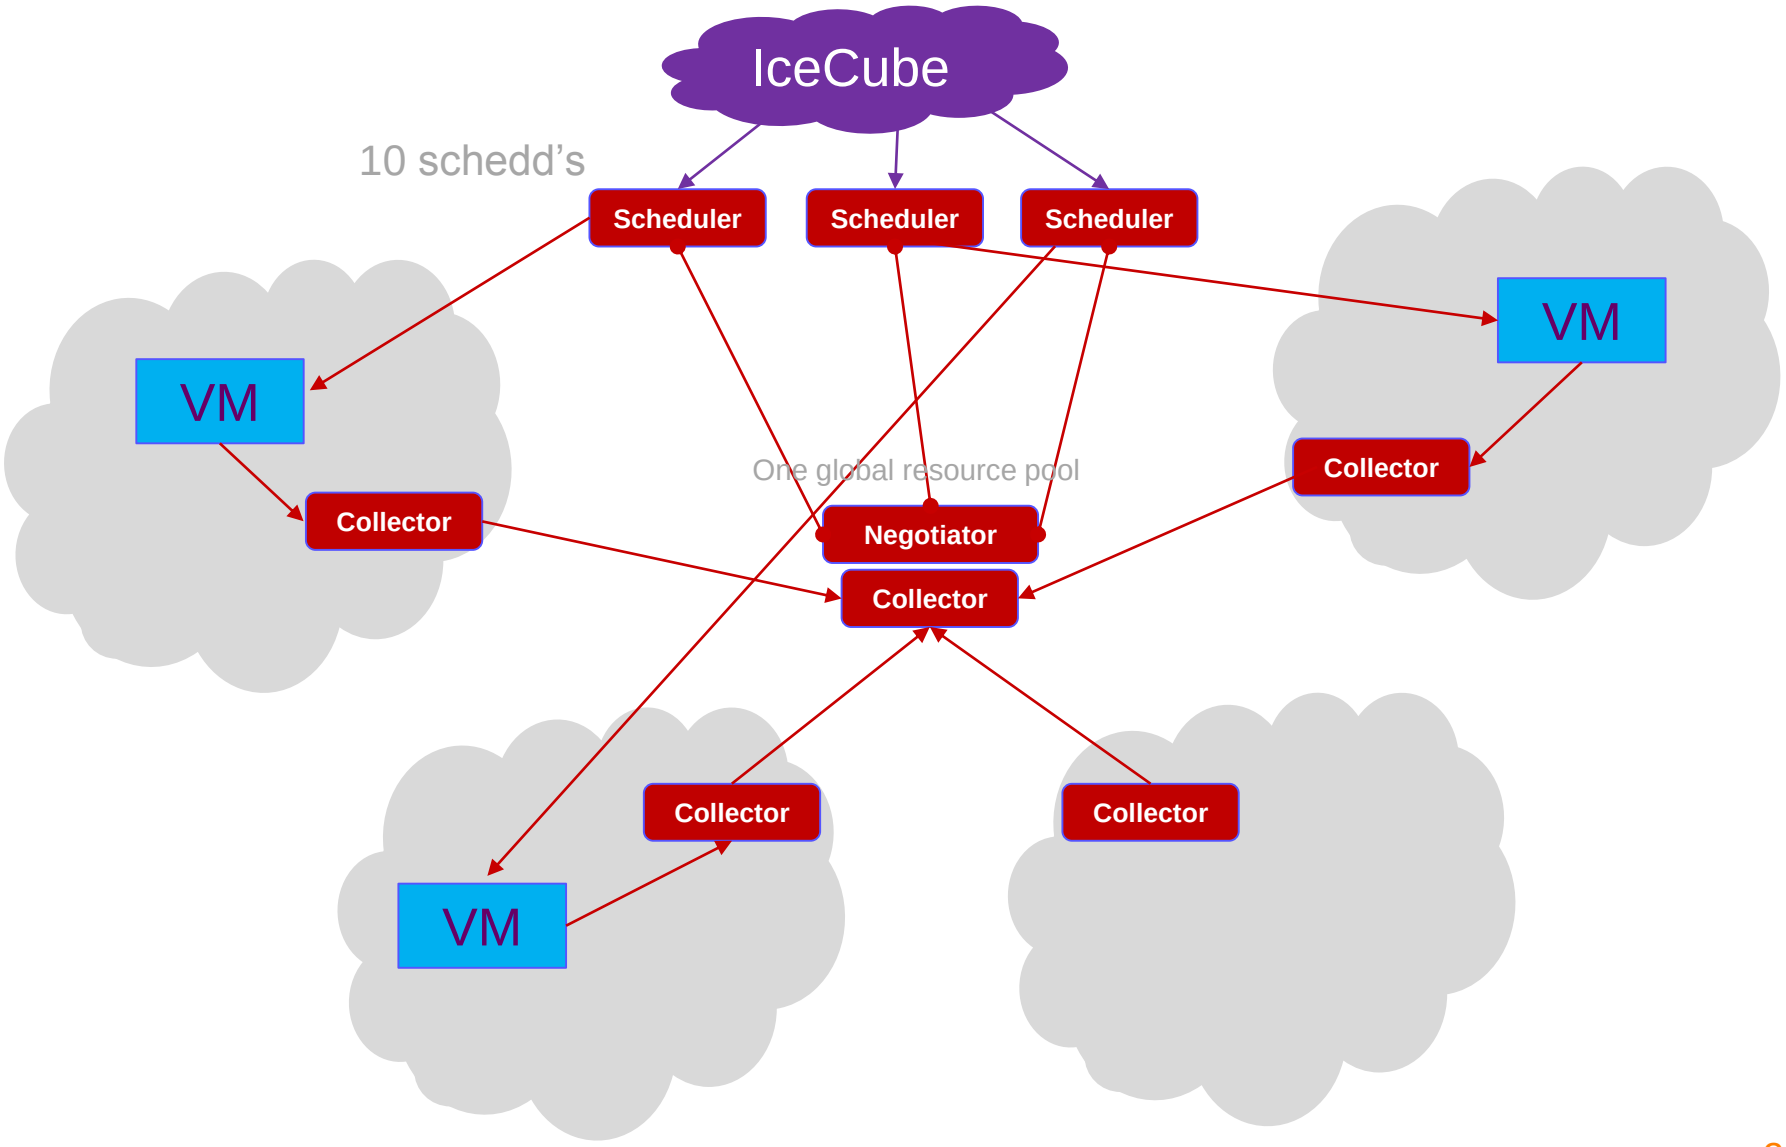

We expect the number of OSG sites to cross the 100 mark by 2021 as a result of NSF-CC* investment in campus clusters (+12 in 2019 and +15 in 2020)

We are pleased to welcome the National Radio Astronomy Observatory (NRAO)

The **European HTCondor Workshop** 2019 was held in Italy, hosted by the European Commission's Joint Research Centre in Ispra, Lombardy, Italy, very close to the shores of Lago Maggiore.

The workshop was the fifth edition in Europe after the successful events at CERN in December 2014, ALBA in February 2016, DESY in June 2017 and RAL in September 2018.

5th Asia Tier Center Forum & 1st Asia HTCondor workshop

24-26 October 2019.
Jointly organized by TIFR Mumbai and KISTI, South Korea
Venue: TIFR, Mumbai India.

Thank you for building such

a wonderful and thriving
(d)HTC community