

Bits & Bytes

No.209, April 2022

Computer Bulletin of the Max Planck Computing and Data Facility (MPCDF)*
<https://docs.mpcdf.mpg.de/bnb>

High-performance Computing

AlphaFold2 on the HPC system *Raven*

AlphaFold2¹ (AF2) is an AI system that is able to predict the 3d structure of a protein from its amino acid sequence. At runtime, the AF2 system initially performs multiple sequence alignments (MSA) on the CPU, followed by the actual structure predictions on the GPU. The MPCDF has been providing installations and job scripts on *Raven* since early August 2021, shortly after Deepmind/Alpha-bet Inc. had publicly released version 2.0.0. This article summarizes recent changes and advancements that are important to run AF2 efficiently on the HPC system *Raven*.

With the deployment of version 2.2.0 (environment module ‘alphafold/2.2.0’) on *Raven*, several improvements were implemented: The AF2 software now runs natively, unlike previous installations that are enclosed in a software container based on the original Docker image. In addition, the job scripts provided by the MPCDF are now split into a plain CPU job for the MSA phase and a dependent GPU job for the subsequent prediction phase. This split helps to minimize idle times of the GPUs. Moreover, to cover the memory footprint of large protein setups it is necessary to allocate more than one GPU together with the host memory via CUDA Unified Memory, as explained in the job scripts. The command

`module help alphafold/2.2.0` gives further instructions to users and information on how to access and use these scripts.

The MSA phase of the AF2 pipeline in particular is I/O bound and puts high load on the file system when reading from the databases. To this end, the MPCDF has stored these databases on a separate file system from the regular user file systems ‘/ptmp’ and ‘/u’ in order to avoid performance impacts on other HPC jobs while maximizing the I/O performance for AF2. In late March 2022, the databases were migrated from spinning disks to a more advanced NVMe-based storage system that is read-only mounted on each *Raven* node. With that change users should experience a significant performance improvement for the MSA phase.

The path to the AF2 databases is provided by the environment variable ‘ALPHAFOLD_DATA’ which is set when an ‘alphafold’ environment module is loaded. Please, only read the databases from the directory referenced via ‘ALPHAFOLD_DATA’, do not create your own copies in ‘/ptmp’ or ‘/u’. Moreover, please do not run AF2 on the HPC system *Cobra* as there is no optimized storage for the databases, and hence the performance is inferior compared to *Raven*.

Klaus Reuter

GitLab CI

GitLab shared runners on GPUs

The MPCDF GitLab instance offers a wide variety of *DevOps* functionalities. One common *DevOps* functionality is *Continuous Integration (CI)*. CI allows the user to define *job pipelines* which are executed after new data was pushed into a GitLab repository. A pipeline can compile code, execute tests or render images, just to name some

common use cases.

These pipelines are executed asynchronously on *GitLab runners*. In principle, every GitLab user can set up a GitLab runner, e.g. on a local laptop or a remote machine the user has access to. In addition, the MPCDF offers several *shared GitLab runners* in the HPC-Cloud² which can be readily employed by any MPCDF GitLab user.

*Editors: Dr. Renate Dohmen & Dr. Markus Rampp, MPCDF

¹<https://www.nature.com/articles/s41586-021-03819-2>

²<https://docs.mpcdf.mpg.de/doc/computing/cloud/index.html>

You can find a list of the currently available runners in the MPCDF documentation³.

Besides the well-established shared runners on CPUs, the MPCDF now offers two additional runners supporting GPU-enabled applications. These runners are labeled *MPCDF-GPU-01* and *MPCDF-GPU02*, and each runner has access to a MIG partition of an Nvidia A30 GPU. They are labeled with the tags *cloud-gpu* and *nvidia-cc80*

and are configured to only run correctly tagged jobs. A CI job that wants to make use of the GPU runners needs to explicitly define at least one of these two tags.

Example: Continuous integration testing of CUDA code

The following example *.gitlab-ci.yml* file demonstrates how to use the GPU shared runners to perform continuous integration testing of a CUDA-enabled HPC code.

```
# .gitlab-ci.yml
cuda-basic-ci:
  image: gitlab-registry.mpcdf.mpg.de/mpcdf/module-image
  tags:
    - cloud-gpu
    - nvidia-cc80
  script:
    - nvidia-smi
    - module avail
    - module load gcc/11 cuda/11.4
    - nvcc --version
  #- ... compile and test the CUDA code as you would do on the HPC system
```

The example uses the ‘module-image’ provided by the MPCDF that offers a software environment largely consistent with the software environments on the HPC systems. As shown the CUDA GPU toolkit and other software are pulled into the environment via `module load` commands.

Continuous integration testing for HPC codes on MPCDF GitLab

Continuous integration (CI) and in particular continuous unit and integration testing are indispensable ingredients for today’s HPC software development workflows. The MPCDF GitLab offers shared runners in the HPC-Cloud that enable teams or individual users to easily perform automated tests for each code commit. Technically, there are shared runners with access to CPUs and GPUs, where each runner offers 4 virtual cores based on the Intel Ice-Lake architecture to the CI jobs, and each GPU-enabled runner offers in addition a virtual GPU corresponding to about 50 % of an A30 GPU (Ampere architecture).

Particularly useful to HPC users is the environment-module enabled Docker image which the MPCDF provides to be used on the shared runners. With that image CI job scripts can simply issue the familiar `module load` commands to get access to virtually the same software as

offered directly on the HPC systems. Hence, tests can easily be implemented for different compilers (e.g. Intel, GNU, Nvidia) or MPI libraries (Intel MPI, OpenMPI). Moreover, within the limits of the 4 virtual cores per runner, sequential vs. parallel execution can be tested based on threads (OpenMP) or processes (MPI), or the correctness of different vectorization levels (e.g. AVX2 or AVX512) can be checked for. Note that these cloud-based resources are less well suited for continuous benchmarking.

The MPCDF recommends to use the shared runners in combination with the ‘module-image’ and not to set up custom individual runners directly on the HPC systems. The latter would not only have to interact properly with the Slurm batch system, but also has an intrinsic security issue for multi-user repositories, because it would execute code committed by any user who can push to the repository in the context of the user who has set up the runner. An example of a ‘gitlab-ci.yml’ file that uses the ‘module-image’ is given in the previous section on the GPU runners. General information about the shared runners is available at the MPCDF documentation pages⁴.

Thomas Zastrow & Klaus Reuter

Globus Online

The MPCDF DataHub service has provided a staging area for multi-terabyte (TB) data transfers for the past several years. In recent times the MPCDF has observed an increasing need for researchers to transfer and share

multi-TB datasets, specifically with non-MPG collaborators from around the world. To address this trend the MPCDF obtained a Globus Online Subscription and on the 9th of March the DataHub’s Globus Online service

³<https://docs.mpcdf.mpg.de/doc/data/gitlab/gitlabrunners.html>

⁴<https://docs.mpcdf.mpg.de/doc/data/gitlab/gitlabrunners.html>

was upgraded to version 5.4 and this subscription was enabled. This upgrade means that extra functionality is now available for MPCDF users via the subscription and that the DataHub now appears in the Globus Online Portal in a slightly different way.

In brief:

- Data is now exposed via collections in the Globus Web Portal (mpcdf#datahub is no longer available).
- Sharing data with any Globus user is now possible.
- Enhanced client functionality is now available to users who join the MPCDF Globus Plus Group.

More detailed information is provided below.

DataHub access via the Globus Online Portal

The upgrade to v5.4 means that the old endpoint “mpcdf#datahub” is no longer available. This was replaced by two new “collections” (logical collections for accessing data). In the Globus Portal these collections are:

1. “MPCDF DataHub Stage-and-Share Area” – The same scratch-based /data area that was mounted on mpcdf#datahub
2. “MPCDF DataHub CBS Project Space” – An explicit collection for the project space of the MPI for Human Cognitive and Brain Sciences

The collections can be found by using the search function in the File Manager or Bookmarks section of the Globus Online Portal.

Accessing the collections remains similar. Simply follow the usual login steps, then link an identity from “MPCDF DataHub OIDC Server (login.datahub.mpcdf.mpg.de)”

and once this is linked use the identity (username@login.datahub.mpcdf.mpg.de) to access the collection.

Enhanced functionality

The new subscription allows the use of the following enhanced functionality:

1. **Data sharing:** In the File Manager section of the Portal a directory can be selected for sharing. This is called a “Guest Collection” and can be shared with individual Globus Online users or groups of users. The guest users can be any user with a Globus Online account, they do not need to have an MPCDF account.
2. **Globus Plus** (for Globus Personal Clients): MPCDF users can now enable sharing from a Globus Connect Personal Endpoint and also perform client-to-client data transfers. This functionality may be enabled by requesting membership of the group “Max Planck Computing and Data Facility Globus Plus”. Simply search for this group in the Groups section of the Portal and click “Join Group” to make a request for membership.

More information

More information can be found in the MPCDF documentation pages: MPCDF DataHub and Globus Online⁵. General information on Globus Online and Globus Connect Personal can be found in the Globus Documentation: How-to⁶, FAQ⁷, Videos⁸. For specific questions about the MPCDF support for Globus Online please create a helpdesk ticket⁹ or mail support@mpcdf.mpg.de.

John Alan Kennedy

New SelfService Features and Improvements

The MPCDF SelfService¹⁰ is constantly evolving. By the end of April version 5.0.0 will be released including additional functionality and improvements of existing workflows to create a better user experience. All changes mentioned below will be made available with this new version.

Redesign of the login process

The login page has been redesigned to be more intuitive and visually appealing. It now simply features an input field for the username and password, respectively, without the need to first specify the type of account. The page

also allows users to request a new MPCDF account via a registration button, following well-established design patterns to provide more clarity.

If two-factor authentication (2FA)¹¹ is activated for your account, the SelfService will ask for the OTP in a second step. In case no OTP can be provided due to a lost or defective token the user can now initiate an automated access restoration workflow without the need to contact the MPCDF support. This will allow users to regain access to their accounts more quickly while maintaining the existing security level. To avoid getting locked out in

⁵<https://docs.mpcdf.mpg.de/doc/data/data-transfer/mpcdf-datahub-and-globus-online.html>

⁶<https://docs.globus.org/how-to/>

⁷<https://docs.globus.org/faq/>

⁸<https://www.globus.org/videos>

⁹<https://helpdesk.mpcdf.mpg.de>

¹⁰<https://selfservice.mpcdf.mpg.de>

¹¹<https://docs.mpcdf.mpg.de/faq/2fa.html>

the first place we strongly recommend creating a backup token¹² to anyone with 2FA enabled.

Viewing accounting data

All users with a regular MPCDF account can now check how many computing and storage resources they used in a given month or time range. This includes computation on our HPC systems as well as storage volume on AFS and other file servers. The data is displayed in multiple tables highlighting different aspects for maximum clarity and control over one's resource usage:

- Ungrouped: this is the raw data that may be used for custom analysis
- By type: see which systems you used most
- By cost center: see how your usage will get billed
- By month: see how your usage changed over time
- By account: in case you own secondary accounts, see which one used how many resources

Institute responsables and accounting departments are able to access the information and tables for their entire institute. The SelfService offers filtering for specific users and/or cost centers.

All tables can be downloaded separately in different formats such as PDF and CSV for documentation or further analysis. Note, however, that the SelfService provides preliminary data for informational purposes only which

may differ from the official accounting and billing.

Additional improvements

The following smaller features will be released with the new SelfService version:

- Users can now see more of their account details as well as their secondary accounts and associated information at "My Account > My data".
- Each account under "My Account > My data" now shows a button "Change password" for easier navigation.
- Supervisors can now filter the list of their supervised users by whether they are locked or not.
- Supervisors can now bulk-edit their supervised users (extend or lock multiple accounts).
- Supervisors can now see the month of the last successful and failed logins for each supervised account. This serves to help supervisors decide whether the account is still needed while keeping the necessary vagueness to disallow user monitoring.

We are aiming to provide a pleasant user experience and are always happy to receive suggestions for improvement and comments on UI from our users. Please send us your comments: support@mpcdf.mpg.de¹³.

Amazigh Zerzour, Andreas Schott

Access to AFS restricted for local Access only

As already announced in Bits&Bytes issue 206¹⁴, the AFS cell *ipp-garching.mpg.de* jointly operated by IPP and MPCDF will be decommissioned in the course of the next few years. As a first step and protective measure the worldwide access to the AFS cell *ipp-garching.mpg.de* will be blocked, while the access will remain possible from the local networks of the IPP (including IPP-HGW), MPA,

MPE, MPQ, MPCDF, and a few external collaboration partners. Also MPP, which is moving to the campus soon, will not be restricted. Respective VPN connections will allow the access, too. This **blocking of AFS-connections** by the firewall will be activated on **May 16th, 2022**.

Andreas Schott

News & Events

AI bootcamp

Together with Nvidia, the MPCDF organizes an online bootcamp "AI for Science", which will take place on May 23rd-24th. The event targets scientists who do not have any prior knowledge of AI methods. During the two-days, hands-on workshop the participants will learn how to apply AI tools, techniques, and algorithms to real-world problems and will study the key concepts of deep neural

networks, how to build deep learning models, and how to assess and improve their accuracy. Since the number of registrations already exceeds the capacity of the workshop, it is planned to organize another "AI for Science" bootcamp towards the end of 2022.

Andreas Marek

¹²<https://docs.mpcdf.mpg.de/faq/2fa.html#how-do-i-enroll-and-use-a-secondary-backup-token>

¹³<mailto:support@mpcdf.mpg.de>

¹⁴<https://docs.mpcdf.mpg.de/bnb/206.html#decommissioning-of-afs>

International HPC Summer School 2022

The International HPC Summer School (IHPCSS) 2022 is planned as an in-person event from June 19th to June 24th in Athens, Greece. The series of these annual events started 2010 in Sicily, Italy. After cancellation of the IHPCSS 2020 event due to the Covid-19 pandemic and conducting IHPCSS 2021 as a pure virtual event, the in-person event in June shall take place with fully vaccinated participants only and under adequate health measures.

For the 2022 event, the organizing partners XSEDE for the US, PRACE for Europe, RIKEN CCS for Japan and the SciNet HPC Consortium for Canada have carried out a joint call for applications. After the reviews by all partners according to the same selection criteria, up to 90 participants have been selected and invited, thereof 30 from European institutions. From Max Planck institutes, four applicants made it into the final selection. School fees, meals and housing will be covered for all accepted applicants. For further information please visit the website of the summer school¹⁵.

Hermann Lederer

Workshop “Introduction to MPCDF services (online)”

The next issue of our semi-annual introductory workshop will be held on April 28th, 14:00-16:30 on zoom. Topics comprise login, file systems, HPC systems, the Slurm batch system, and the MPCDF services remote visualization, Jupyter notebooks and datashare. Basic knowledge of Linux is required. Registration is necessary and can be done here¹⁶.

Tilman Dannert

“Meet MPCDF”: New online forum and lectures for MPCDF users

In June 2022, the MPCDF will launch a new series of monthly online lectures together with a Q&A forum for its users. The event will be held on the first Thursday of the

month, from 15:30 to 16:30, and features a technical talk (ca. 20-30 minutes) about an AI, HPC or data-related topic given by a staff member of the MPCDF. In addition, this meeting offers the opportunity for the users to informally interact with MPCDF staff, in order to discuss relevant kinds of technical topics. Optionally, questions or requests for specific topics to be covered in more depth can be raised in advance via e-mail¹⁷. Target audience are intermediate to advanced users of MPCDF services as well as computational scientists and software developers of the MPG. Users seeking a basic introduction to MPCDF services instead, are referred to our semi-annual online workshop “Introduction to MPCDF services” (see above) which offers a consistent introduction for new users of the MPCDF.

The first “Meet MPCDF” event will take place on June 2nd at 15:30 with the talk “Introduction to the AI tools at the MPCDF”. The second issue is planned for July 7th and will cover a topic on HPC software engineering. Connection details and updates can be found on the MPCDF webpage¹⁸.

Tilman Dannert, Markus Rammpp

RDA-Deutschland-Tagung 2022

February 21st-25th, more than 500 people attended this year’s online conference of the German chapter¹⁹ of the Research Data Alliance²⁰. Various topics around research data management were discussed with a specific focus on the FAIR²¹ (*Findable, Accessible, Interoperable and Reuseable*) as well as the CARE principles²² – *Collective Benefit, Authority to control, Responsibility and Ethics*. The schedule and some of the slides presented are available from the conference website²³. As in previous years the MPCDF was involved in the organization of the conference. Next year’s conference is scheduled for February 13th-17th, 2023.

Raphael Ritz

¹⁵<https://ss22.ihpcss.org/>

¹⁶<https://events.gwdg.de/e/intro-mpcdf-2022-04-28>

¹⁷<mailto:training@mpcdf.mpg.de>

¹⁸<https://www.mpcdf.mpg.de/services/training>

¹⁹<https://www.rda-deutschland.de/>

²⁰<https://www.rd-alliance.org/>

²¹<https://force11.org/info/the-fair-data-principles/>

²²<https://www.gida-global.org/care>

²³<https://indico.desy.de/event/31380/timetable/#20220221>