

CDO — Climate Data Operators

Uwe Schulzweida, Oliver Heidmann, Ralf Müller², and Luis Kornblueh

uwe.schulzweida@mpimet.mpg.de, Max Planck Institute for Meteorology
² ralf.mueller@dkrz.de, Deutsches Klimarechenzentrum

What is CDO ?

CDO is a collection of many operators for standard processing of climate and forecast model data.

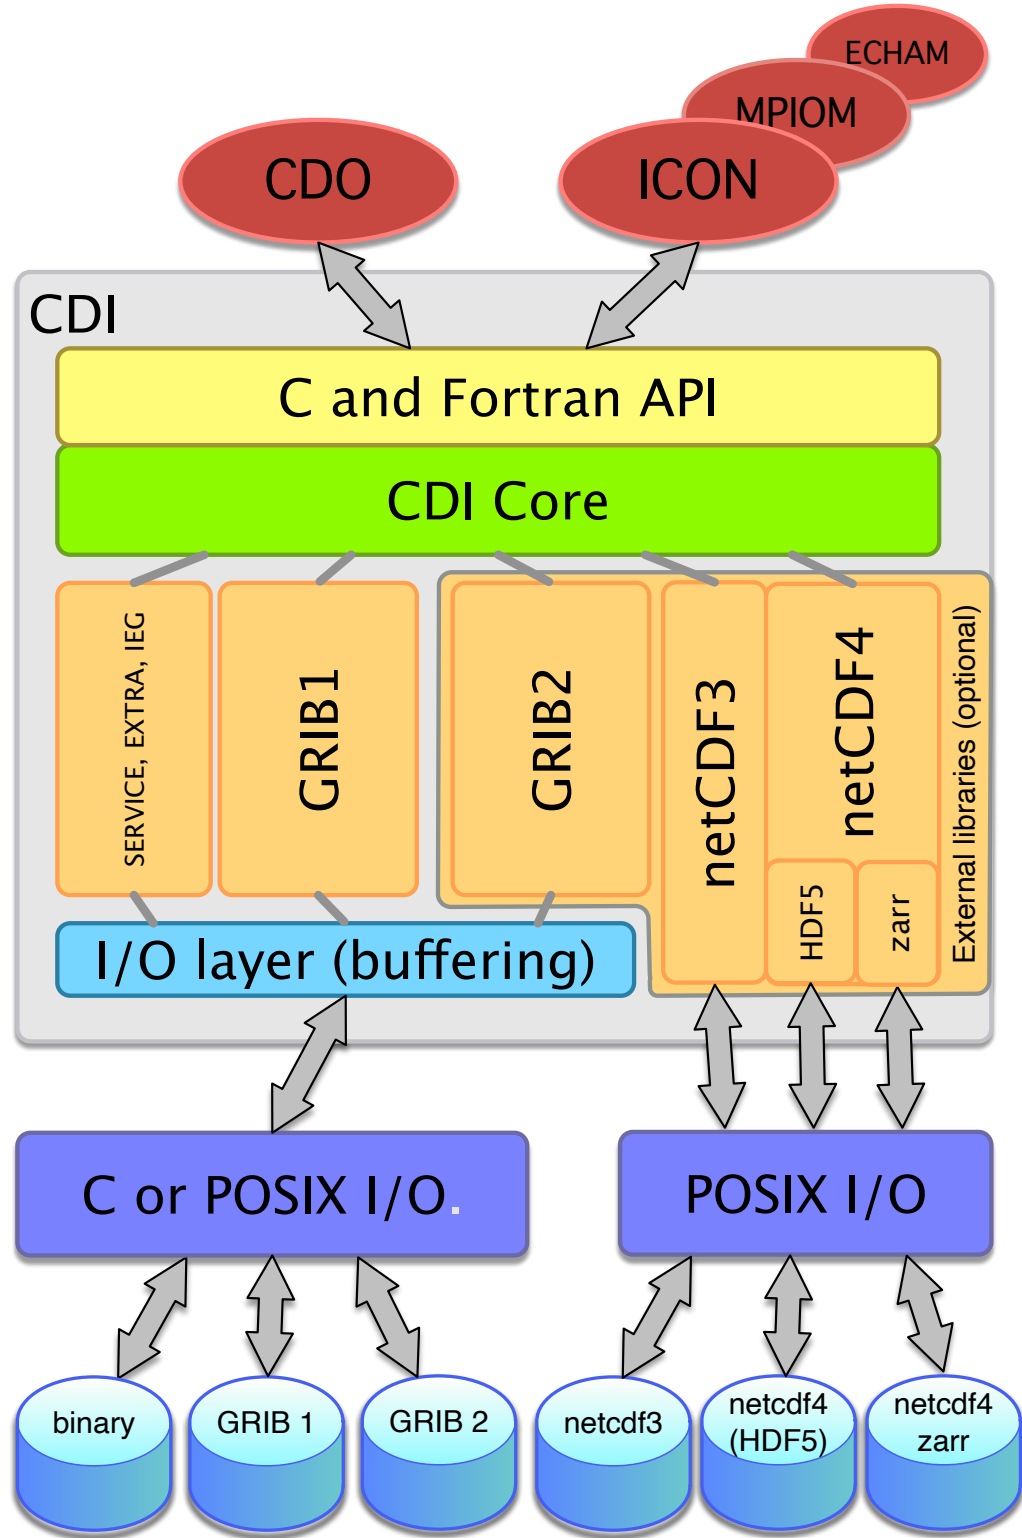
- File format conversion: GRIB ↔ netCDF
- Interpolation between different grid types and resolutions
- Portability (ISO C++17 with some POSIX extensions)
- Performance (fast processing of large datasets, multi-threaded)
- Modular design and easily extendable with new operators
- UNIX command line interface
- Script-language interface for Python and Ruby
- Tested on many UNIX/Linux systems, Cygwin, and MacOS-X

Data I/O Interface

CDO is using the I/O interface CDI (Climate Data Interface) which is shared with all major MPI-M climate models (for models with asynchronous, parallel I/O).

- GRIB1 via CGRIBEX (MPI-M)
- GRIB2 via ECCODES (ECMWF)
- netCDF, CF-convention (UNIDATA)
- SERVICE, EXTRA, IEG (MPI-M legacy binary formats)

GRIB support includes highly efficient, fast compression algorithms.




Available Operators

CDO provides more than 600 operators which can be pipelined on thread level. CPU time intensive operators are OpenMP parallelized.

File information	Print information about datasets
File operations	Copy, split and merge datasets
Selection	Select parts of a dataset
Comparison	Compare datasets
Modification	Modify data and metadata
Arithmetic	Arithmetically process datasets
Statistical values	Ensemble, spatial and temporal statistic
Interpolation	Spatial and temporal interpolation
Import/Export	HDF5, binary, ASCII
Climate indices	ETCCDI and ECA Indices
Plot utility	based on Magics++, ECMWF

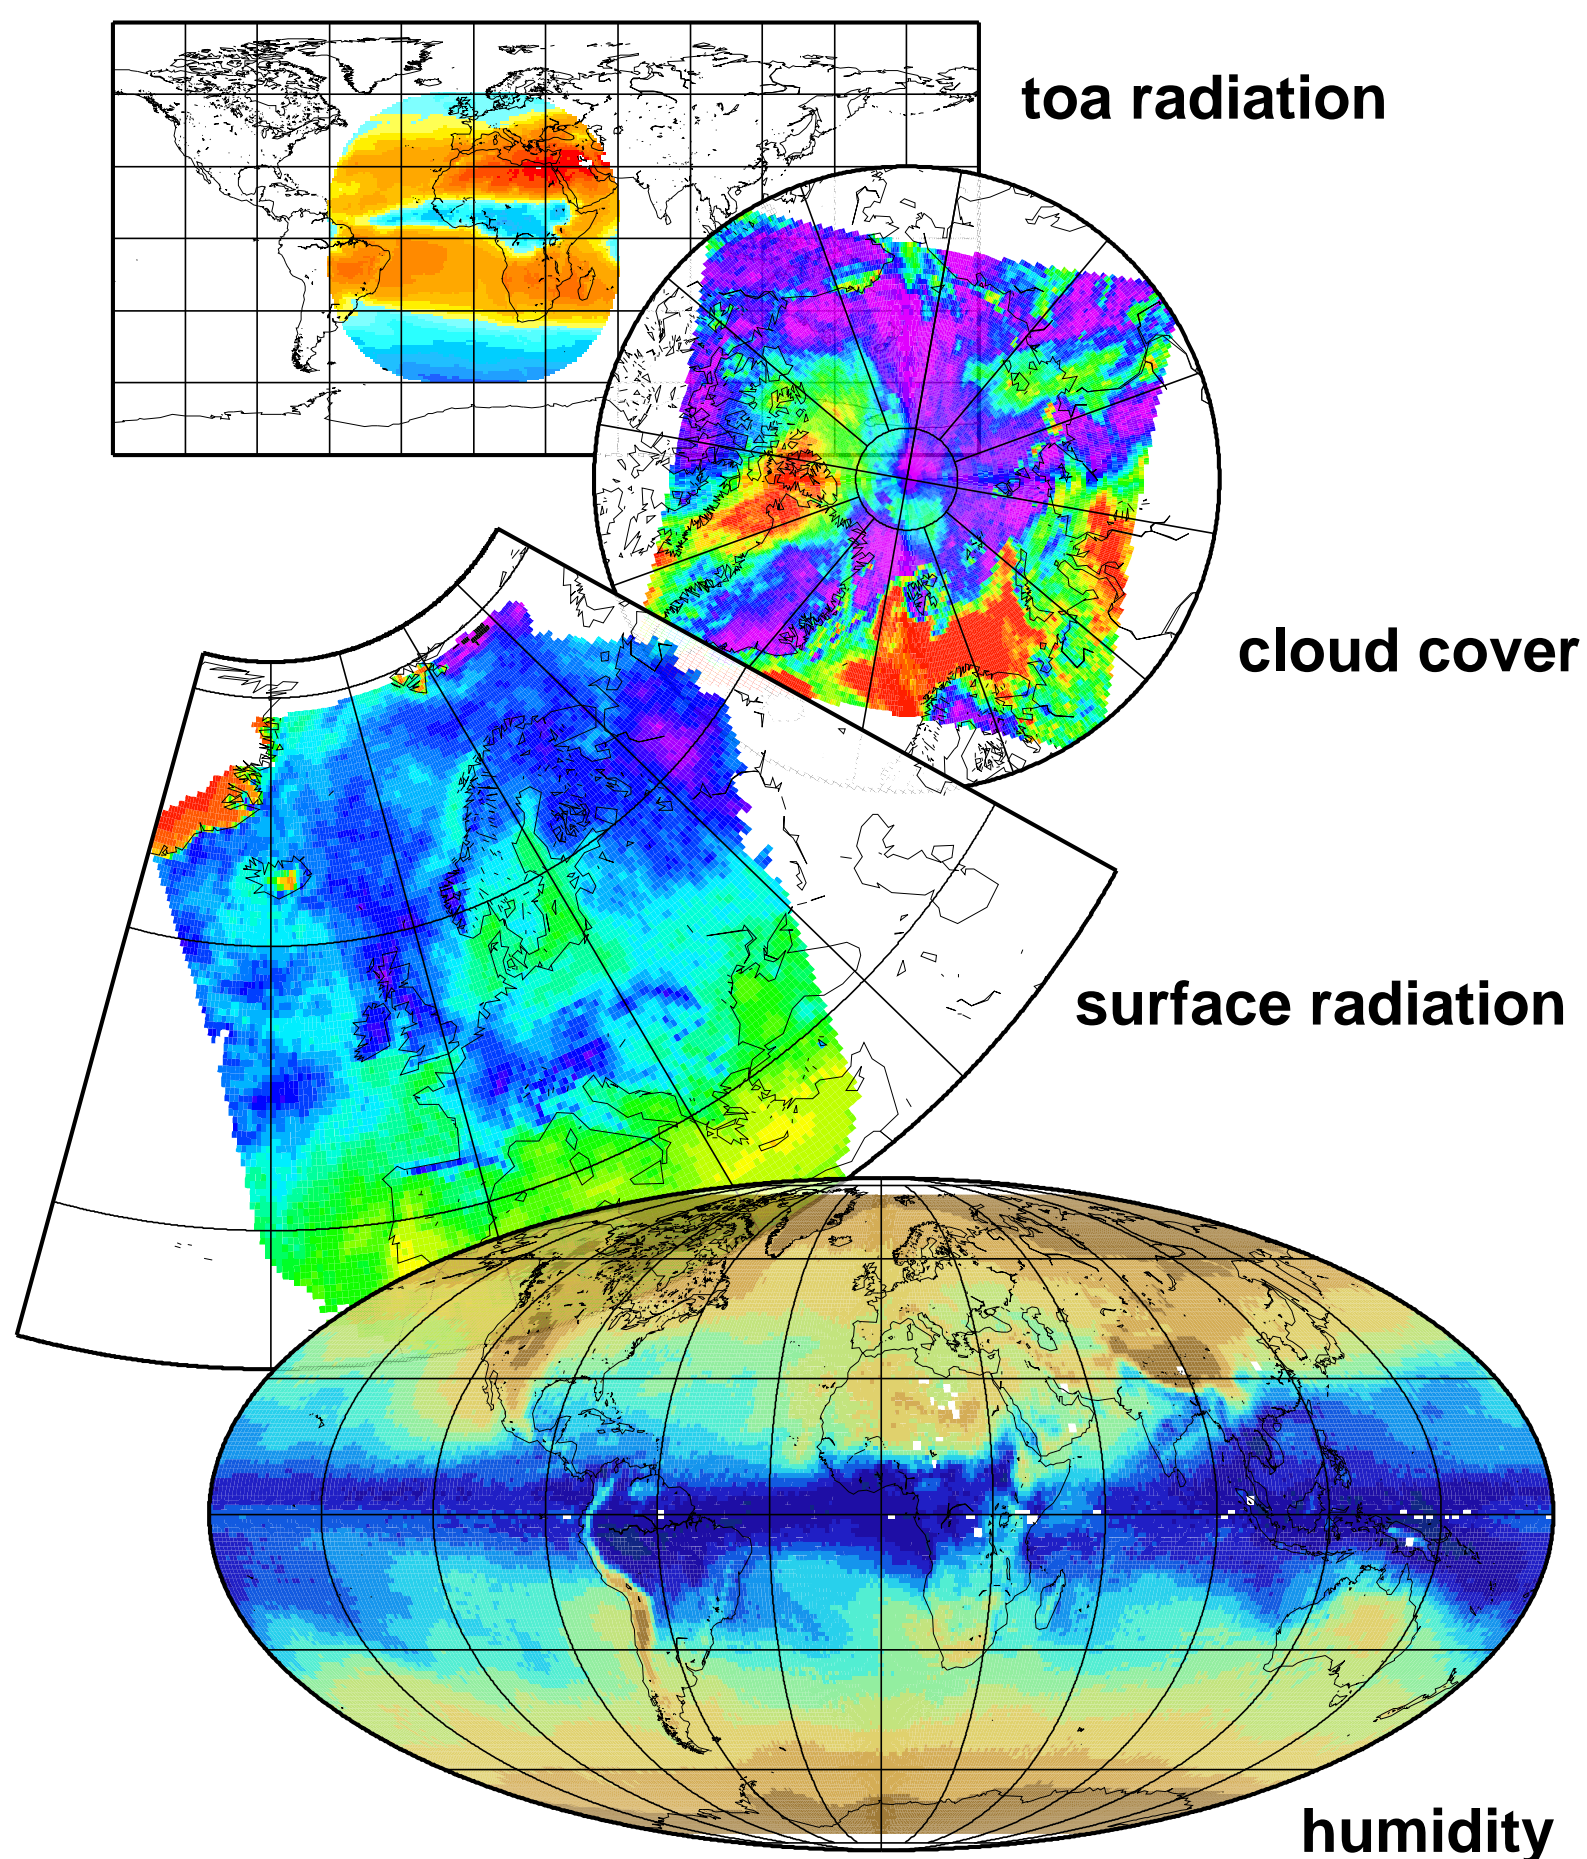
Supported Grids

A large set of grids is supported including spectral- and Fourier-coefficients. HEALPix, Gaussian, regular and rotated lat-lon grids, conformal mapped quadrilateral grids, and finally general unstructured grids.



Satellite-data Support

EUMETSAT's Climate Monitoring Satellite Application Facility provides satellite-derived geophysical parameters for climate monitoring. Data sets contain several cloud parameters, surface albedo, radiation fluxes, temperature and humidity profiles. These products are stored in HDF5. DWD has funded the import CDO operator import_cmsaf.

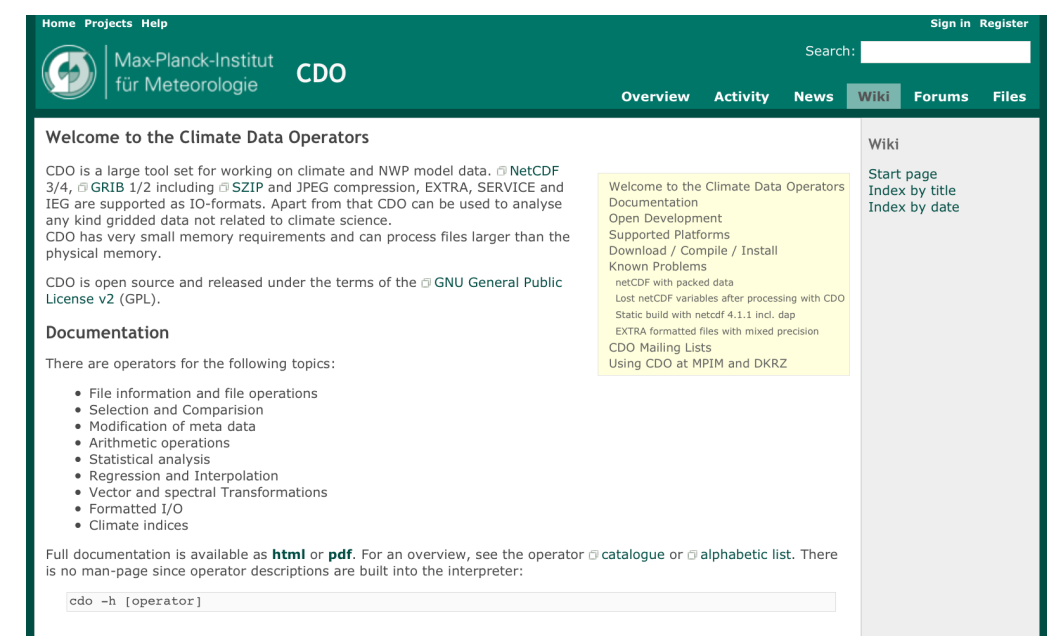


Usage Examples

- Converts a GRIB formatted dataset to netCDF4:
`cdo -f nc4 copy ifile ofile`
- Mean over 3 ensemble members:
`cdo ensmean ifile1 ifile2 ifile3 ofile`
- Expressions – vector for wind speed from u-v components:
`cdo expr, 'wind speed=sqrt(u*u+v*v)' ifile ofile`
- Import a CM-SAF product in HDF5 format, interpolate to a global one degree grid, and convert the result to netCDF:
`cdo -f nc remapbil,global 1 -import cmsaf ifile.h5 ofile.nc`

Community Support

A fully featured development platform is available to support the community at <https://code.mpimet.mpg.de/projects/cdo>.



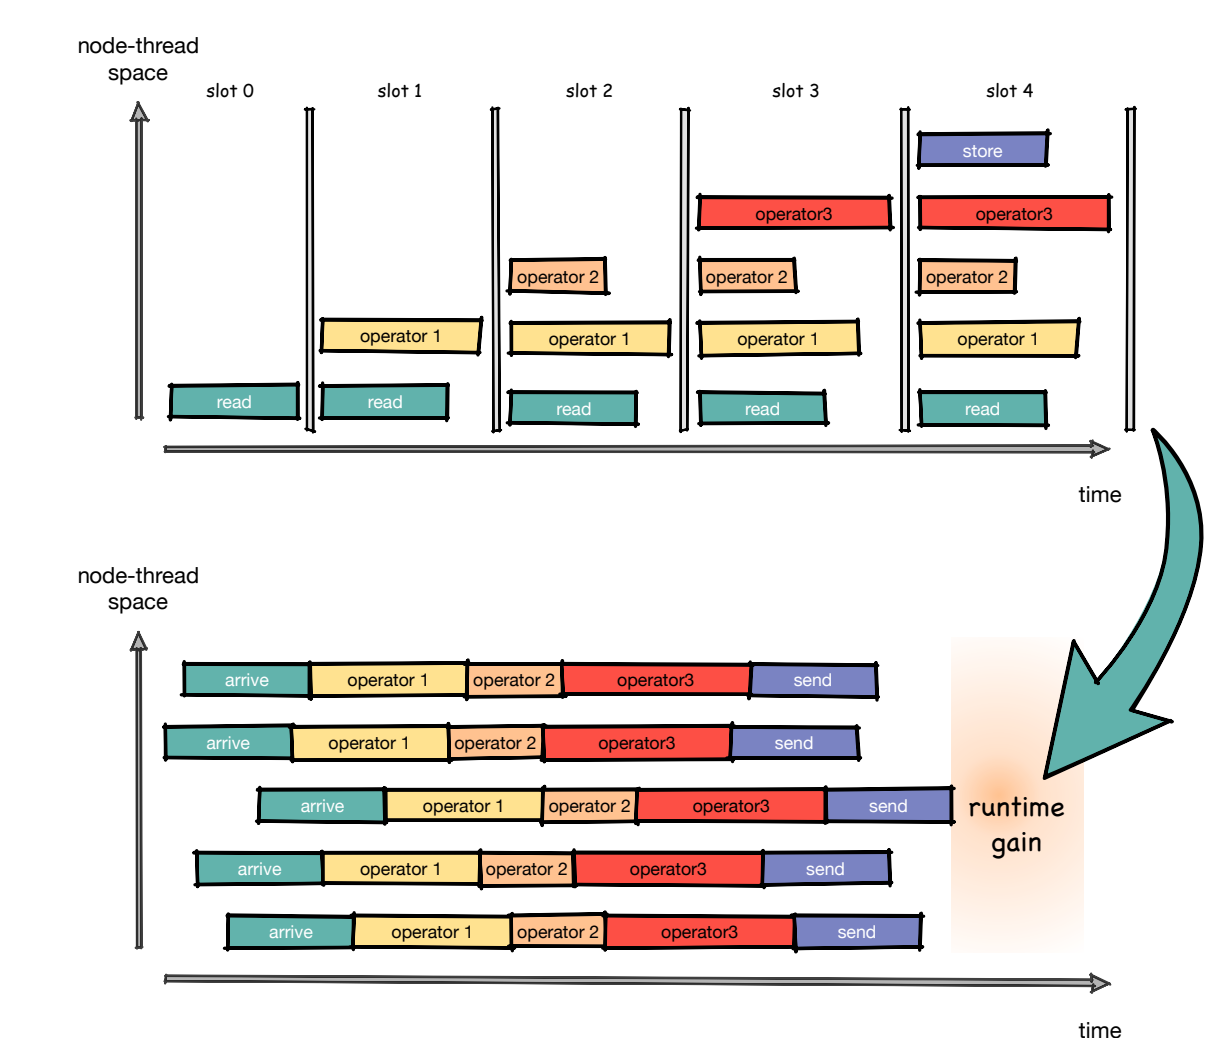
- User wiki
- Documentation
- Bug tracking system
- User forums
- Download area
- Repository access
- Sponsored by IS-ENES

Latest features added

- Asynchronous decompression of GRIB records
- HEALPix grid support
- Zarr data support via NCZarr
- Bit rounding support
- Improved expression evaluation operators

What is coming next?

- Librarization of CDO
- Asynchronous direct acyclic graph based processing of data (optimization of subtasks within CDO)



Runtime is improving by exploring fine-grained parallelism: it is necessary to change from a pipelined to a scheduler task processing model.