

# Towards performance benchmarks for VT26 on Virgo

A. Redelbach, S. Schütt

16/07/2026

FIAS, Goethe Universität

# Setup on VIRGO

- Input: vt26test.00001.tsa
- Script: run\_cbmreco\_1.sh
- Compilation of Cbmroot within vae26.hpc.gsi.de (thanks Pierre-Alain)
- Timing information: Adding -t option for cbmreco binary (thanks Bartosz)
- First testing, variation of OpenMP threads

## Slurm settings:

- cpus-per-task=1 to 256
- reserve (on work partition)
- Minimal settings first, could be optimized

# Results for CPU scaling on VIRGO

- Machines ccexe0132, ccexe0142, ccexe0145
- Significant acceleration „reserved“ vs. „testing mode“ (more than factor of 2)
- Other times can also be extracted, first focus on CA tracking

| Num. of OpenMP threads | CA runtime (in s) |
|------------------------|-------------------|
| 1                      | 370 (no res.)     |
| 2                      | 190               |
| 4                      | 98                |
| 8                      | 51                |
| 16                     | 23                |
| 32                     | 11                |
| 64                     | 4                 |
| 128                    | 1.5               |
| 256                    | 1.1               |

# Approach for GPU utilization on VIRGO

- Details of setup (event-based or time-based) to be defined
- Input from Grigory:
  - `sbatch --account=cbm --export=ALL, PARS_DIR=/lustre/cbm/users/.../parameters, IMAGE=/lustre/cbm/users/.../containers/cbm_online-master.sif, OUTPUT_DIR=/lustre/cbm/users/.../output submit_cbmreco_hydra_mi100.sbatch`
  - sbatch script including aptainer with binding
  - `XPU_DEVICE=hip0`
  - option: `ca-gpu-multiwindow-length-ns`
- cbm\_online-master as standard online container → can be pulled to a lx machine
- Need to be included in the `video` group on amd\_gpu nodes
- Request sent for access to `hydra.hpc.gsi.de` for submitting GPU jobs

General point: Where to display updates (Mattermost, Indico, Redmine)?