

Multi-Order HEALPix Map implementation in CDS HEALPix Rust

François-Xavier Pineau¹

¹Centre de Données astronomiques de Strasbourg

IVOA, College Park, 06 June 2025



□ Introduction

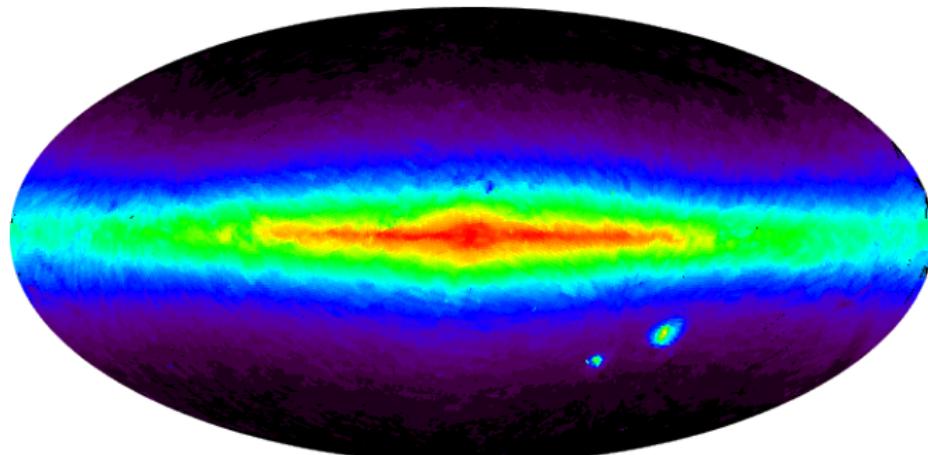
The HEALPix ecosystem encompass various “maps” (often leading to confusion).

By chronological order:

- **Skymap** (flat resolution HEALPix maps)
 - FITS convention from the official HEALPix lib
- **MOC: Multi-Order Coverage map**
 - IVOA standard, v1.0 serialisation inspired from skymaps
- **BMOC: Binary Multi-Order Coverage map**
 - limited to CDS HEALPix (Java and Rust)
- **MOM: Multi-Order HEALpix Map**
 - Martinez-Castellanos, Singer et al. 2021, used in Ligo/Virgo
 - FITS serialisation also inspired from skymaps

□ Skymaps (flat resolution)

- HEALPix maps, aka *Skymaps*
- Description:
 - list of **cells at a same resolution** (regular hpx grid)
 - a quantity/**value is associated to each cell**
- Example:
 - Ligo/Virgo flat probability map
 - 2MASS count/density map



MOM in CDS HEALPix

□ Skymaps (flat resolution)

- FITS serialisation (several conventions)
 - convention from the official HEALPix library
 - doc. asked by M. Taylor (TOPCAT), according to P. Fernique
 - EXPLICIT maps: BINTABLE with IPIX + values

FITS file format for HEALPix products
Version 0.2

E. Hivon on behalf of the HEALPix collaboration

June 29, 2018

Abstract

This document describes the requirements on the FITS files containing HEALPix products, and in particular sky maps

Contents

1 Introduction	1
2 Formats	1
2.1 Sky Maps	1
2.1.1 Full sky maps	2
2.1.2 Cut sky maps	2

This page describes data format conventions for FITS-blinned data and model representations provided with the HEALPix algorithm.

HEALPix Formats

This section describes a proposal for HEALPix format conventions which is based on formats currently used within the Fermi Science Tools (ST) and pointlike. This format is intended for representing maps and cubes of both integral and differential quantities including:

- Photon count maps and cubes (e.g. as generated with ghein).
- Exposure cubes (e.g. as generated with gexpcube2).
- Source maps - product of exposure with instrument response in spatial dimension (e.g. as generated with gchesmap).
- Model maps and cubes (the Fermi IEM and other diffuse-emission components).

The format defines a **SKYMAP HDU** for storing a sequence of image slices (bands) and a **BANDS HDU** to store the geometry and coordinate mapping for each band. A band can represent any selection on non-spatial coordinates such as energy, time, or FoV/angle. The most common use-case is a sequence of bands representing energy bins (for counts map) or energy nodes (for source or model maps).

<https://gamma-astro-data-formats.readthedocs.io/en/latest/skymaps/healpix/>

□ Skymaps (flat resolution)

- Well supported:
 - official HEALPix libraries (HEALPy), TOPCAT, Aladin, ...
 - partly supported in CDS HEALPix/MOC Rust/Python (read in MOCPy; write IMPLICIT in CDS HEALPix Python)

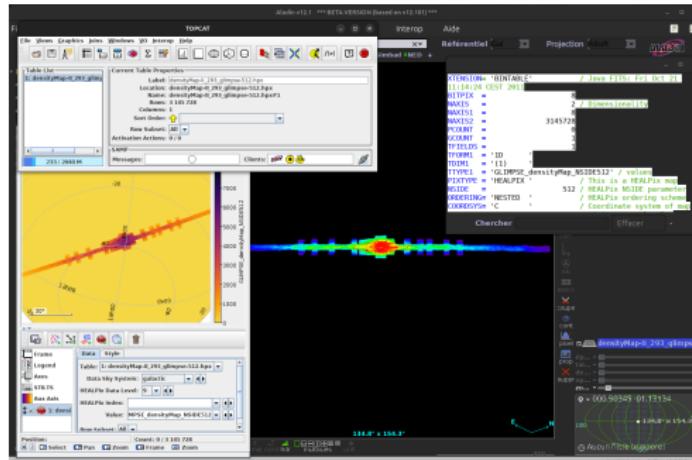
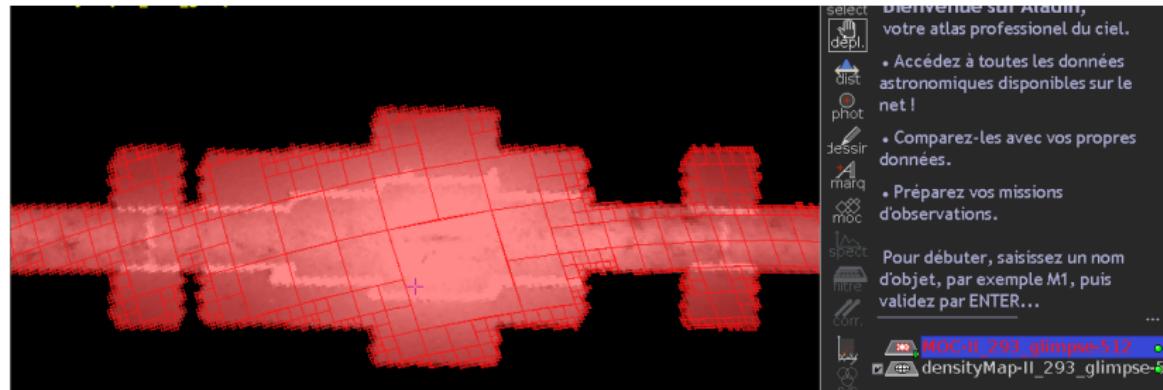


Figure 1: Load Glimpse hpx file (T. Boch) in TOPCAT and Aladin

MOC

- MOC: Multi Order Coverage map
- Description:
 - list of **non-overlapping cells at various resolutions**
 - **4 siblings cells must be merged** into the parent cell
 - a MOC is a coverage/a mask on the unit-sphere
- Example
 - coverage of the glimpse survey



MOC

- MOC is an IVOA standard
- FITS serialization
 - v1.0 inspired from skymaps + uses NUNIQ
 - NUNIQ: ORDER + IPIX packed together
 - v2.0 serialisation relies on ORDER max (29) ranges
 - faster set operations (not/and/or/xor)
 - no NUNIQ for Time (and Frequency) qualities
- For implementations/tool support see IVOA MOC page
 - Java:
 - Official/CDS library used in Aladin (P. Fernique, M. Reinecke), TOPCAT, ...
 - Official HEALPix lib (v2.0 support?)
 - Python:
 - [MOCPy](#) (Rust based)
 - [PyMOC](#) (pure python), no v2.0 support (yet?)
 - Rust + WebAssembly: CDS MOC Rust (+moc-cli), MOCWasm
 - ...

□ BMOC

- BMOC: **Binary Multi-Order Coverage map**
- Description:
 - list of **non-overlapping cells at various resolutions**
 - a **binary value** (0 or 1) is **associated to each cell**
 - 0: the cell is partially covered; 1: the cell is fully covered
 - 4 siblings cells **having a same flag value** are merged into the parent cell
 - “special” set operations (not/and/or/xor)
- Example
 - see Fast (possibly complex) STC-S queries thanks to B-MOCs

□ BMOC

- No document, no paper
 - only implemented in CDS HEALPix Java / Rust
 - **New!** command line tool `hpx-cli`
- Currently 2 FITS serializations (**new!**)
 - both rely on the ZUNIQ numbering, similar to `google s2geometry`
 - one custom: BIGENDIAN in primary HDU optimize FITS size and allows for memmap
 - one interoperable: BINTABLE with ZUNIQ and flag
- BMOC from the command line with `hpx-cli`

```
> hpx cov 9 stcs "Circle ICRS 147.6 69.9 0.4"
> hpx cov -t fits -o bmoc.fits \
  9 stcs "Circle ICRS 147.6 69.9 0.4"
> hpx cov -t bintable -o bmoc.fits \
  9 stcs "Circle ICRS 147.6 69.9 0.4"
```

MOM

- MOM: Multi Order skyMaps
- Description:
 - list of **non-overlapping cells at various resolutions**
 - a **quantity/value is associated to each cell**
- Example:

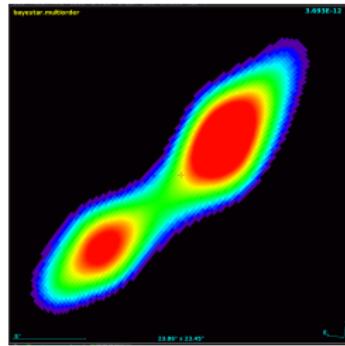


Figure 2: Ligo/Virgo Multi-Order probability Map

MOM

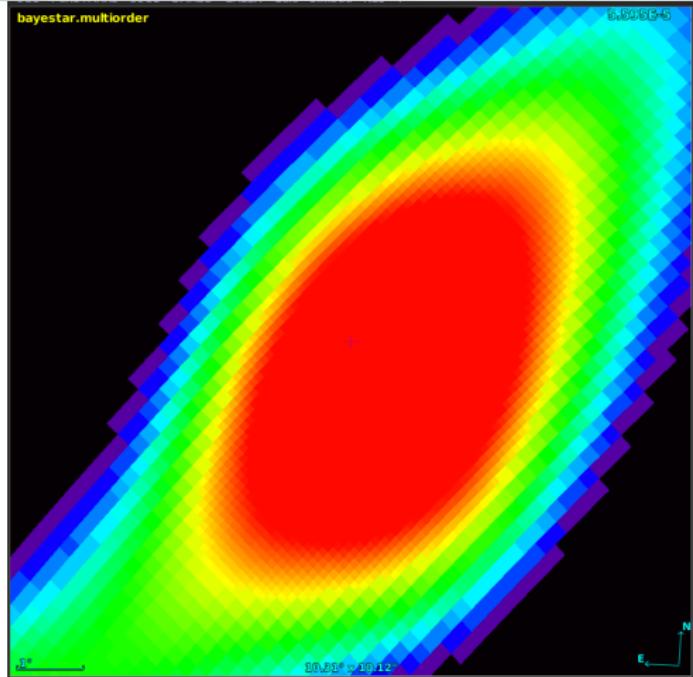


Figure 3: Ligo/Virgo Multi-Order probability Map

MOM

- Origin
 - Martinez-Castellanos, Singer et al. 2021: “*Multi-Resolution HEALPix Maps for Multi-Wavelength and Multi-Messenger Astronomy*”
 - C++ implementation in `mhealpy`
 - Now, **Multi-Order Sky Maps**
- Similar idea at CDS: “*Valued-MOC*”, now **MOM**
 - BMOC (<2018) in CDS HEALPix (Java and Rust)
 - particular MOM with value = boolean
 - single bit added to a UNIQ like numbering
 - operations dedicated to boolean case, no genericity
 - MOM implementation in CDS HEALPix Rust (since 09/2024)
 - very generic, very different from `mhealpy`?
- Other examples
 - HATS structure: merge siblings while `#src < 1M`
 - see also Rust/Python code by K. Malanchev
 - **New!** χ^2 MOM
 - **New!** x-match spurious association maps

□ MOMs in CDS HEALPix Rust

- MOM creation from a skymap is very generic
 - type of index is generic (u32, u64, ...)
 - type of value is generic (f32, i64, list, ...)
 - function merging (or not) siblings is generic
 - takes parent cell order and ipix
 - takes the 4 siblings values
 - return either the merged value or None
- single requirement from the skymap: iterate on (ipix, value) tuples on ascending ipix

```
fn from_skymap_ref<'s, S, M>(skymap: &'s S, merger: M) -> Self
where
    S: SkyMap<'s, HashType = Self::ZUniqHType, ValueType = Self::ValueType>,
    M: Fn(u8, Self::ZUniqHType, [&Self::ValueType; 4]) -> Option<Self::ValueType>,
    Self::ValueType: 's;
```

□ χ^2 MOMs

- Build a χ^2 MOM from a density/count Skymap
- Example with [hpx-cli](#)
- Input file `gaia_edr3_dist.csv`
 - 1 467 744 819 rows
 - 10 columns
 - 165 GB

```
> head -10 gaia_edr3_dist.csv
source_id,RA_ICRS,DE_ICRS,r_med_geo,r_lo_geo, ...
1000000057322000000,104.87975539563,55.95486459306,784.5591
1000000121746472704,104.85627575018,55.96825004364,1875.965
1000000156106220032,104.86295428908,55.97882280840,2708.931
1000000156106221440,104.85273193584,55.98099099003,1336.845
...
...
```

□ χ^2 MOMs

- Density map FITS file creation

```
> time cut -f , -d 2,3 gaia_edr3_dist.csv \
| tail -n +2 \
| hpx map dens 11 gaia_edr3_dist.dens.fits list -d ,
```

```
real    8m25,812s (=> 340 MB/s or 2.9e+6 src/s)
user    11m7,242s
sys     3m40,669s
```

- PNG creation from the 385 MB FITS file

```
> time hpx map view --silent gaia_edr3_dist.dens.fits \
gaia_edr3_dist.dens.png allsky 400
```

```
real    0m20,361s (14s on my laptop)
user    0m19,246s
sys     0m0,492s
```

\square χ^2 MOMs

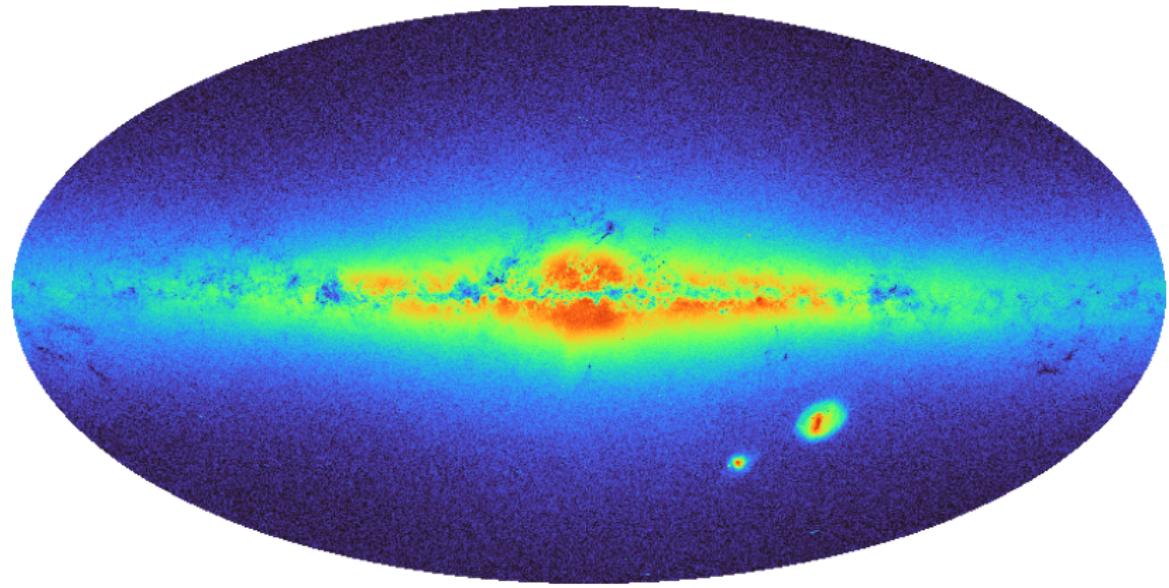


Figure 4: Skymap at order 11

□ χ^2 MOMs

- Create a χ^2 MOM from the Skymap
 - **from 385 MB to 13 MB in 1s**

```
> time hpx map convert gaia_edr3_dist.dens.fits \
   gaia_edr3_dist.dens.mom.fits dens2chi2mom
```

```
real    0m0,978s
user    0m0,661s
sys     0m0,317s
```

- PNG creation from the 13 MB FITS file

```
> time hpx mom view --silent gaia_edr3_dist.dens.mom.fits \
   gaia_edr3_dist.dens.mom.png allsky 400
```

```
real    0m0,494s
user    0m0,478s
sys     0m0,013s
```

χ^2 MOMs

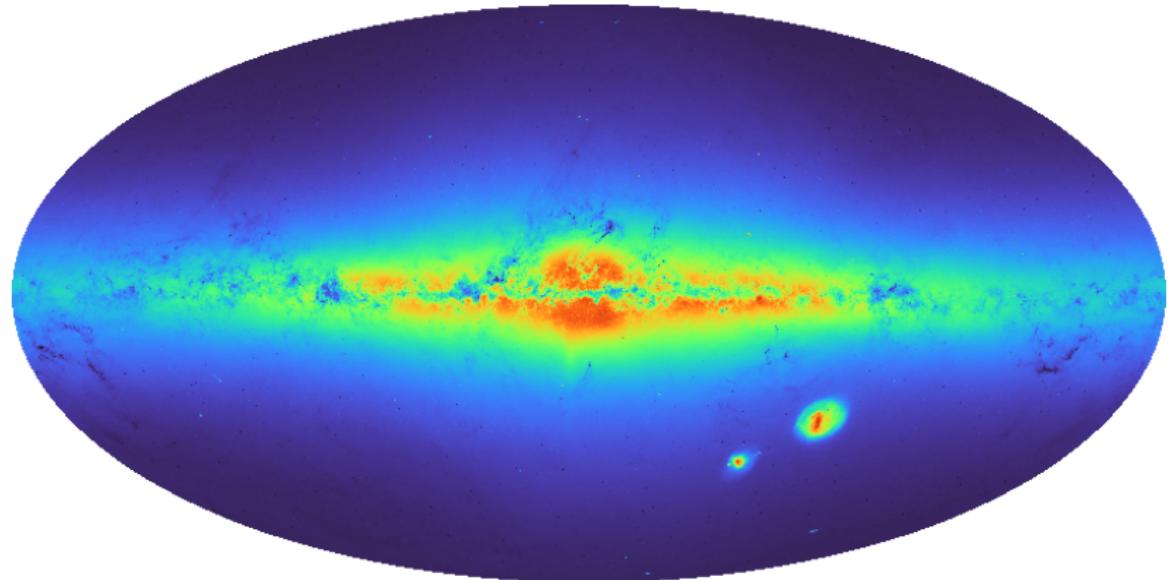


Figure 5: χ^2 MOM

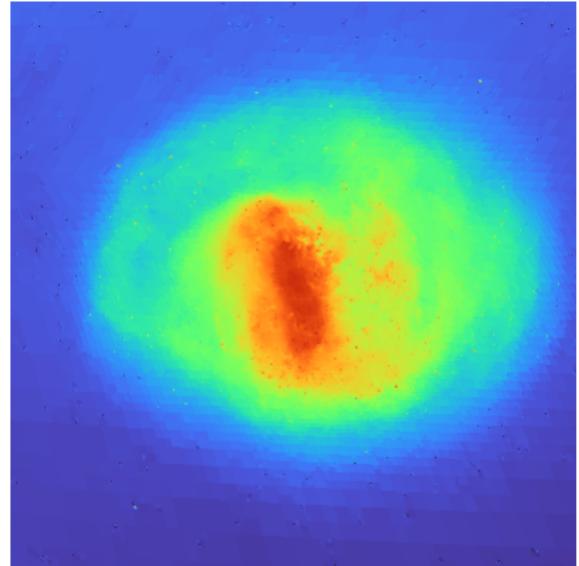
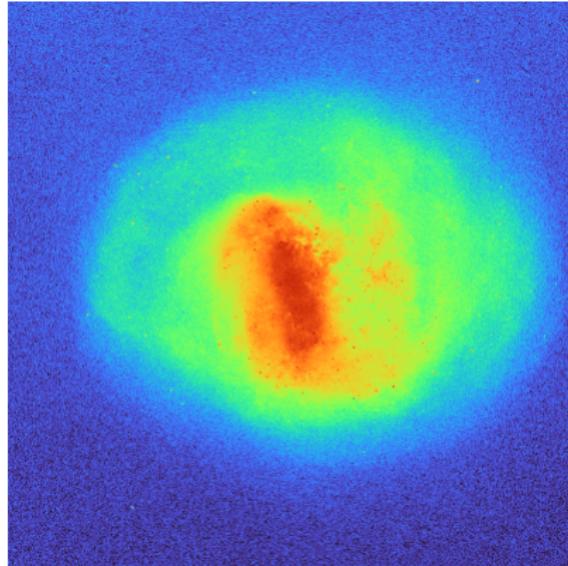
□ χ^2 MOMs

Zoom around the LMC

```
> hpx map view --silent gaia_edr3_dist.dens.fits \
  gaia_edr3_dist.dens.lmc.png \
  custom -l 280.4652 -b -32.8884 \
  sin 400 400 \
  --x-bounds [-0.15..0.15] --y-bounds [-0.15..0.15]

> hpx mom view --silent gaia_edr3_dist.dens.mom.fits
  gaia_edr3_dist.dens.mom.lmc.png \
  custom -l 280.4652 -b -32.8884 \
  sin 400 400 \
  --x-bounds [-0.15..0.15] --y-bounds [-0.15..0.15]
```

\square χ^2 MOMs



Left: skymap. Right: χ^2 MOM (we see quantization due to merged cells in low density regions; horizontal alignment due to galactic projection)

□ χ^2 MOMs

- On the previous example:
 - 385 vs 13 MB: **x30** compression
 - PNG creation 20 vs 0.5 s: **x40** faster
 - White noise removed in χ^2 MOM!
 - While keeping high resolution where density varies quickly (order 11 $\sim= 1.7\text{arcmin}$)

□ HATS struct MOM

Instructions:

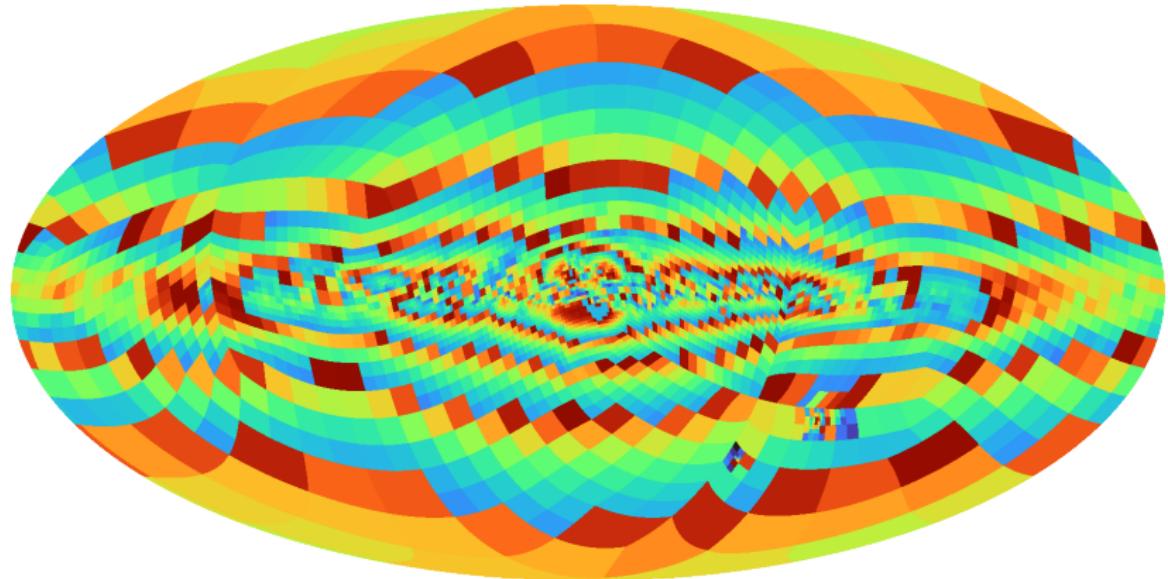
- download [Skymap file](#) (T. Boch) at order 10 ($3.4' \times 3.4'$)
- create a MOM from the Skymap with hpx-cli
- convert the FITS MOM into FITS BINTABLE with hpx-cli
- load in TOPCAT
- view with hpx-cli

□ HATS struct MOM

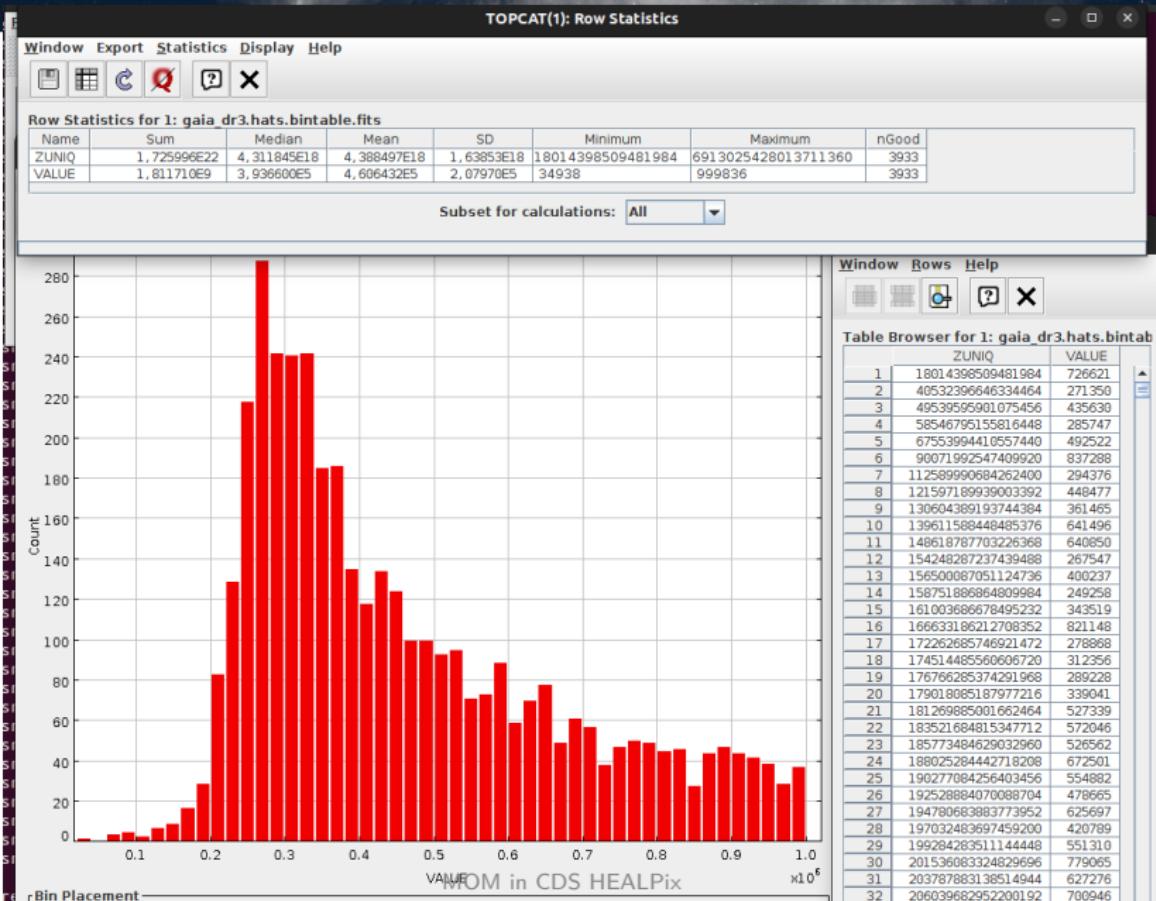
Command lines:

```
> wget -O gdr3.d10.skymap.fits "${URL}"  
> hpx map convert gdr3.d10.skymap.fits gdr3.hats.fits \  
    count2mom 1000000  
> hpx mom convert gdr3.hats.fits gdr3.hats.bintable.fits \  
    bintable  
> topcat gdr3.hats.bintable.fits &  
> hpx mom view gdr3.hats.fits gdr3.hats.png \  
    -c linear allsky 800
```

□ HATS struct MOM



HATS struct MOM



□ MOMs in CDS HEALPix Rust

- Operation between 2 MOMs also **very generic**:
 - requires two MOMs of same Keys and Values types
 - a split function to split a parent cell into 4 siblings
 - an operation returning an optional value from either:
 - one value from the left MOM
 - one value from the right MOM
 - one value from both MOMs
 - a possible post-merge operation

```
fn merge<'s, L, R, S, O, M>(lhs: L, rhs: R, split: S, op: O, merge: M) -> Self
where
    L: Mom<'s, ZUniqHType = Self::ZUniqHType, ValueType = Self::ValueType>,
    R: Mom<'s, ZUniqHType = Self::ZUniqHType, ValueType = Self::ValueType>,
    S: Fn(u8, Self::ZUniqHType, Self::ValueType) -> [Self::ValueType; 4],
    O: Fn(LhsRhsBoth<Self::ValueType>) -> Option<Self::ValueType>,
    M: Fn(
        u8,
        Self::ZUniqHType,
        [Self::ValueType; 4],
    ) -> Result<Self::ValueType, [Self::ValueType; 4]>;
```

□ Operation between 2 MOMs

- Examples
 - merge 2 probability MOMs
 - implemented in Ligo/Virgo (`mhealpy`) I guess
 - not implemented (yet?) in CDS HEALPix Rust
 - just requires to code split, op and merge
 - rate of spurious x-match
 - Density MOM 1 times density MOM 2 time constant value (cone-search area)
 - post-merge χ^2 compatible number of spurious matches

```
pineau@cds-dev-fxp:~$ hpx mom op mult-and-chi2-merge --help
Multiply left values by right values by a constant and post-chi2 merge. Values must be f32 or f64 and are assumed to be densities (i.e. remain the same splitting a cell)

Usage: hpx mom op <LHS_FILE> <RHS_FILE> <OUT_FILE> mult-and-chi2-merge [OPTIONS]

Options:
  -c, --cte <CTE>          Constant to be used in the multiplication [default: 1]
  -n, --value-name <NAME>    Name assigned to the computed value in the output file
  [default: MULT]
  -t, --threshold <THRESHOLD> Completeness of the chi2 distribution of 3 degrees of
  freedom below which cells are post-merged [default: 16.266]
  -h, --help                  Print help
```



The question of the representation/serialization.

□ MOC Representation

Evolution of the MOC internal representation and FITS serialisation:

- MOC 1.0: **by layer** representation
 - Java library: one IPIX array per ORDER
 - FITS: sorted list of NUNIQ
 - low resolution cells before high resolution cells
- MOC 2.0: **max order range** representation
 - both in memory and in FITS
 - **faster** set operations (not, and, or, ...)
 - operation possible on **streammed** MOCs
 - **No NUNIQ for Time (v2.0) and Frequency (v2.1)**
- v2.0 may be more compact than v1.0, not always
 - x2 in favor of v1.0 in case of isolated cells

□ MOC representation

NUNIQ verus SUNIQ

- NUNIQ: described in the **HEALPix paper**
 - base cell $\in [0, 11]$ coded on 4 bits ($[0, 15]$)
 - NUNIQ: add 4 to the base cell number
 - no extra bit required
 - $NUNIQ = (4 \ll (ORDER \ll 1)) + IPIX$
- **But**, no extra value available in bit patterns of Time or Frequency cells
- SUNIQ: based on **sentinel bit**
 - Marks the limit of the highest possible value
 - SUNIQ add 16 to the base cell number
 - requires an extra bit
 - $SUNIQ = (16 \ll (ORDER \ll 1)) | IPIX$
- SUNIQ solves the problem of new MOC quantities!
- SUNIQ natural ordering still based first on ORDER, then IPIX

□ MOC representation

- For easy and efficient (and stream) operations on MOC, BMOC, MOM we **need another sorting scheme**.
- **MOC 2.0** representation based on max resolution ranges
 - cells **ordering does not depends on cell ORDERs**
 - (follows the Z-Order curve ordering)

□ MOM representation

- MOM
 - each cell has a distinct value thus **successive cells cannot be merged into ranges**
 - **efficient/stream operation: need cells to follow the Z-Order curve!!**
- How to add the ORDER information in the IPIX, preserving the max res ordering?
 - **sentinel bit on the LSB part instead of on the MSB part!**
 - adopted by [google s2geometry](#)
 - $ZUNIQ = ((IPIX \ll 1)|1) \ll ((ORDMAX - ORDER) \ll 1)$

□ Unified view

- Today:
 - MOC v1.0: sorted list of NUNIQ
 - MOC v2.0: sorted list of MAX ORDER ranges
 - BMOC: sorted list of **ZUNIQ** with extra bit flag
 - MOM:
 - Ligo/Virgo: sorted list of NUNIQ, VALUES tuples
 - CDS HEALPix: sorted list of **ZUNIQ**, VALUE tuples
- But:
 - NUNIQ valid only for HEALPIX
 - MOM could be generalized to Time and Frequency
 - NUNIQ not adapted for
 - streaming operations (large MOMs)
 - simple and efficient operations
- Use **ZUNIQ** everywhere?!

□ Open questions

- Description in a standard/note of SUNIQ and/or ZUNIQ?
- MOC is ia particular BMOC, BMOC is a particular MOM (Multi Order SkyMap)
 - similarities but various serializations
 - should we try to homogenize them?
- One proposition:
 - drop NUNIQ: no streamming, not compatible with Time, Frequency, ...
 - remark: 1D count MOM = adaptative histogram
 - use ZUNIQ:
 - from/to ZUNIQ comp. time = from/to ZUNIQ compt. time
 - natural ordering allows for **stream processing/fast operations**
- Any ideas/interests?

□ Wrap-up side

- MOM recently implemented in CDS HEALPix Rust
- Very generic implementation
 - core algo implemented
 - custom needs: implement basic functions
- A few use case examples:
 - merge probability MOMs
 - χ^2 density MOMs
 - compute x-match probabilities
 - ... any use case from you!
- FITS serialisation
 - mHEALPy: NUNIQ (internal representation??)
 - CDS HEALPix: ZUNIQ
- I favor ZUNIQ for MOMs, but
 - also for BMOCs (with a separated flag)?
 - also for MOCs?