

Radio-astronomical imaging with HIP

Stefano Corda

stefano.corda@epfl.ch

04-10-2022

Swiss SKA Days, Lugano

In collaboration with:

Bram Veenboer

ASTRON





Square Kilometre Array Challenge

SKA (Square Kilometre Array) requirements (per Science Data Processor (SDP) site):


- ~ 260 PFlop/s
- ~ 157 TB/s
- ~ 5 MW

SKA1 LOW - the SKA's low-frequency instrument

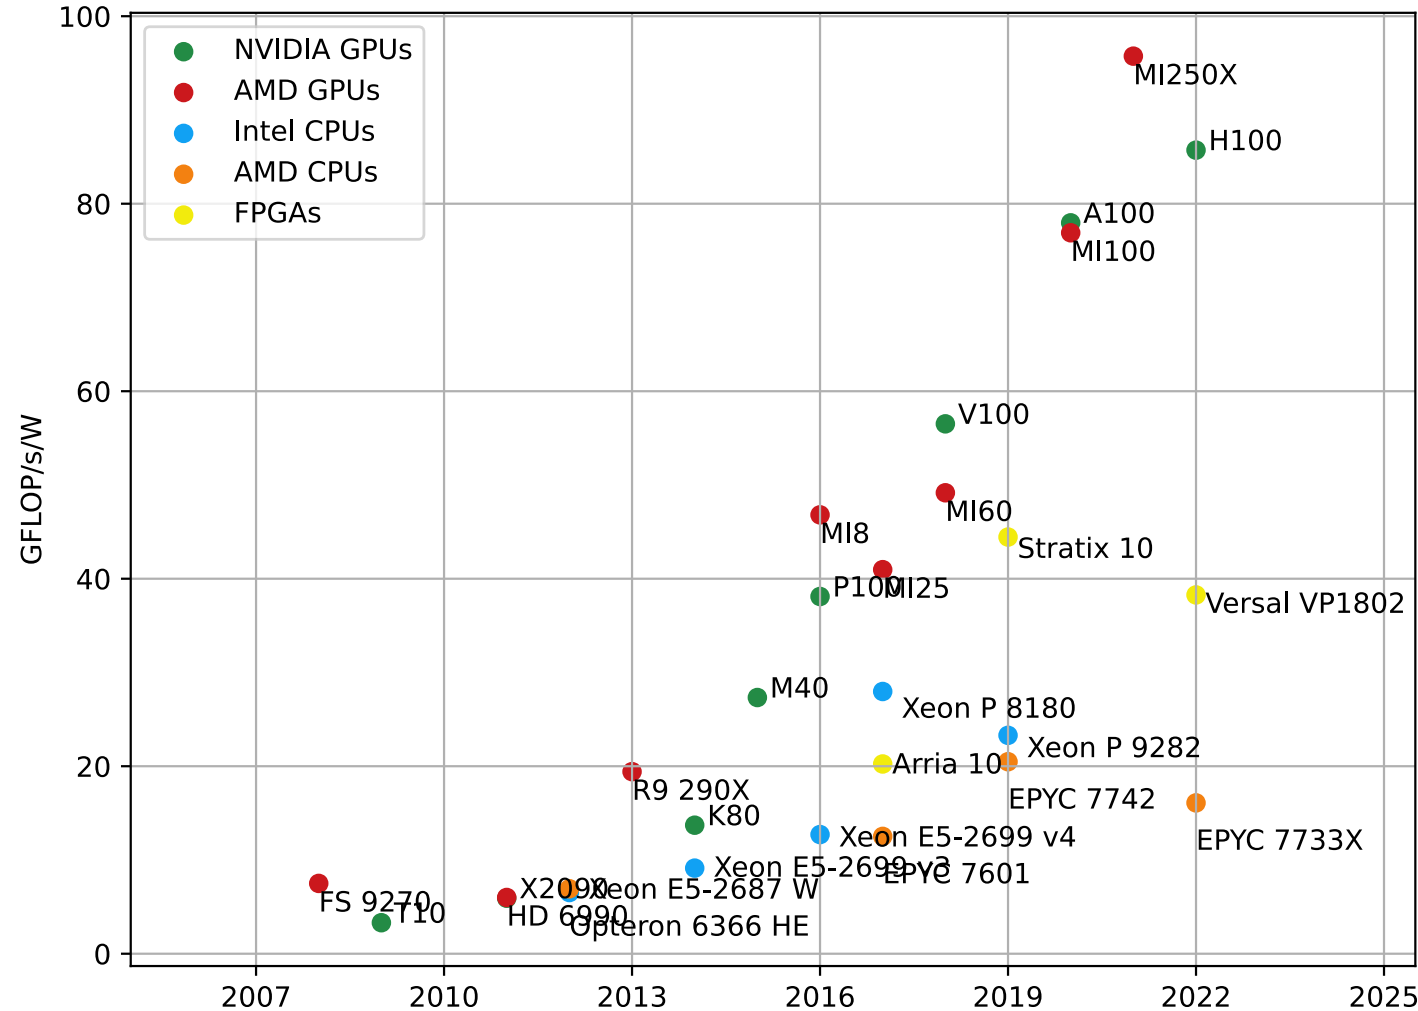
The Square Kilometre Array (SKA) will be the world's largest radio telescope, revolutionising our understanding of the Universe. The SKA will be built in two phases - SKA1 and SKA2 - starting in 2018, with SKA1 representing a fraction of the full SKA. SKA1 will include two instruments - SKA1 MID and SKA1 LOW - observing the Universe at different frequencies.




Hardware trend




Hardware trend



- Accelerators (GPUs and FPGAs) achieve better performance and energy efficiency than CPUs.
- Interesting competition in the GPU market (AMD vs NVIDIA).
- AMD systems are in the first positions of TOP500 and Green500 (ISC 22).

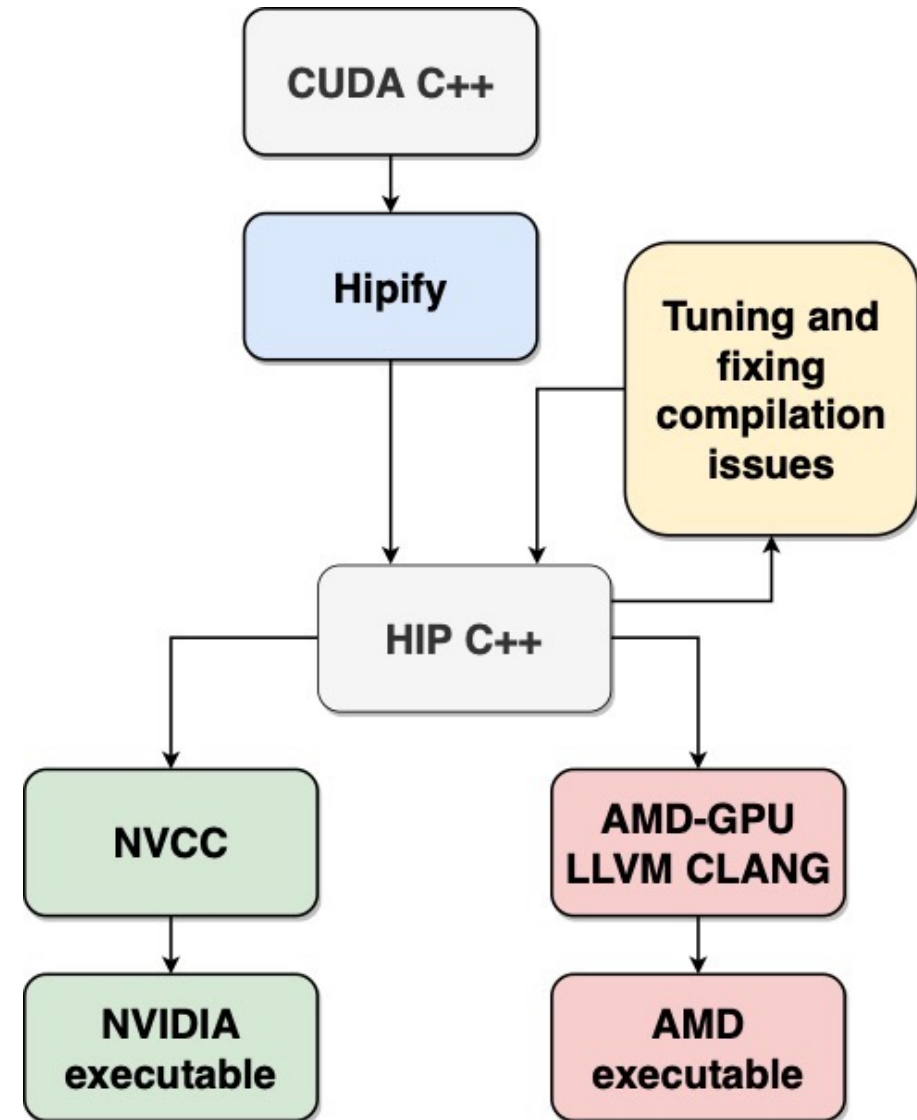


CPU




GPU

HIP



HIP

Portability: code is compiled with NVIDIA or AMD compilers.

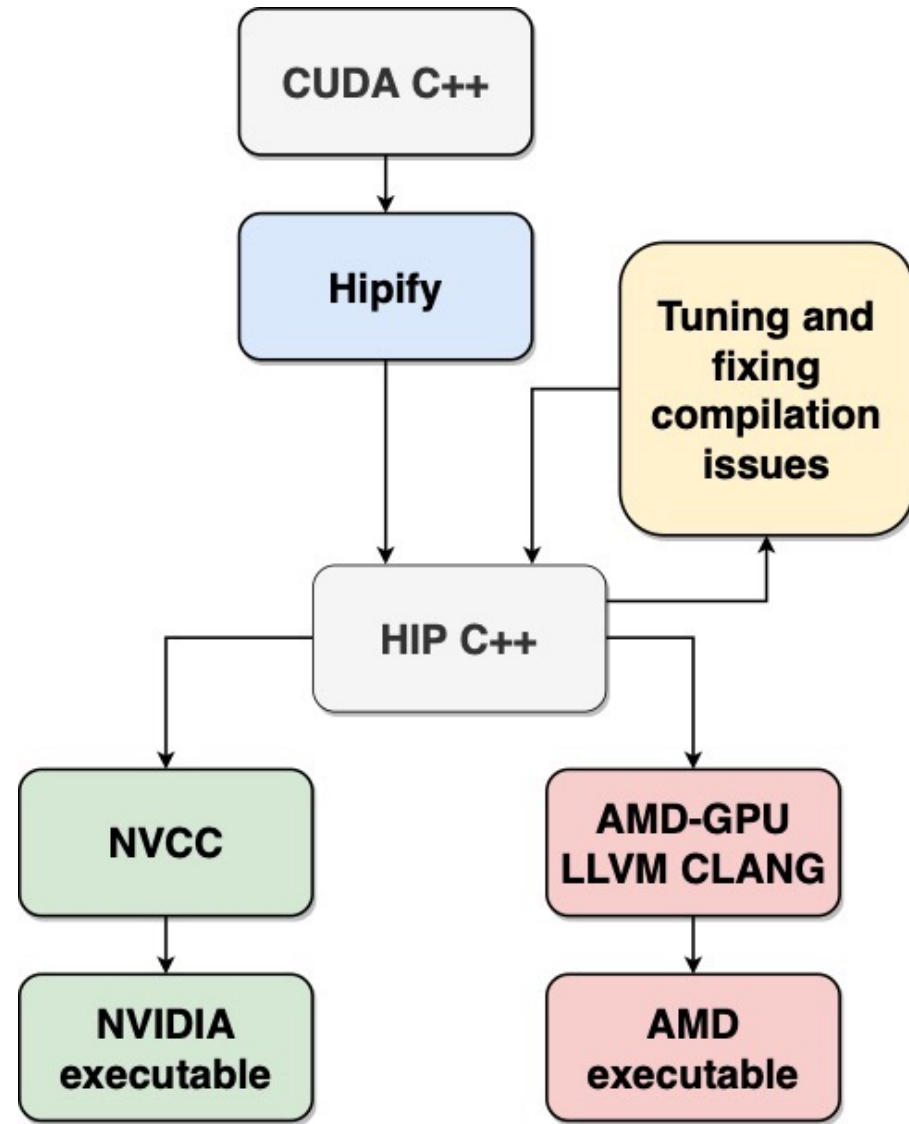


HIP

Portability: code is compiled with NVIDIA or AMD compilers.

Programmability: semantic is similar to CUDA (> 99%).

- First attempt (industry partner from France):
- Porting *all* CUDA code to HIP with Hipify
- Fixed some compilation issues
- Same performance on NVIDIA GPUs
- Poor performance on AMD GPUs



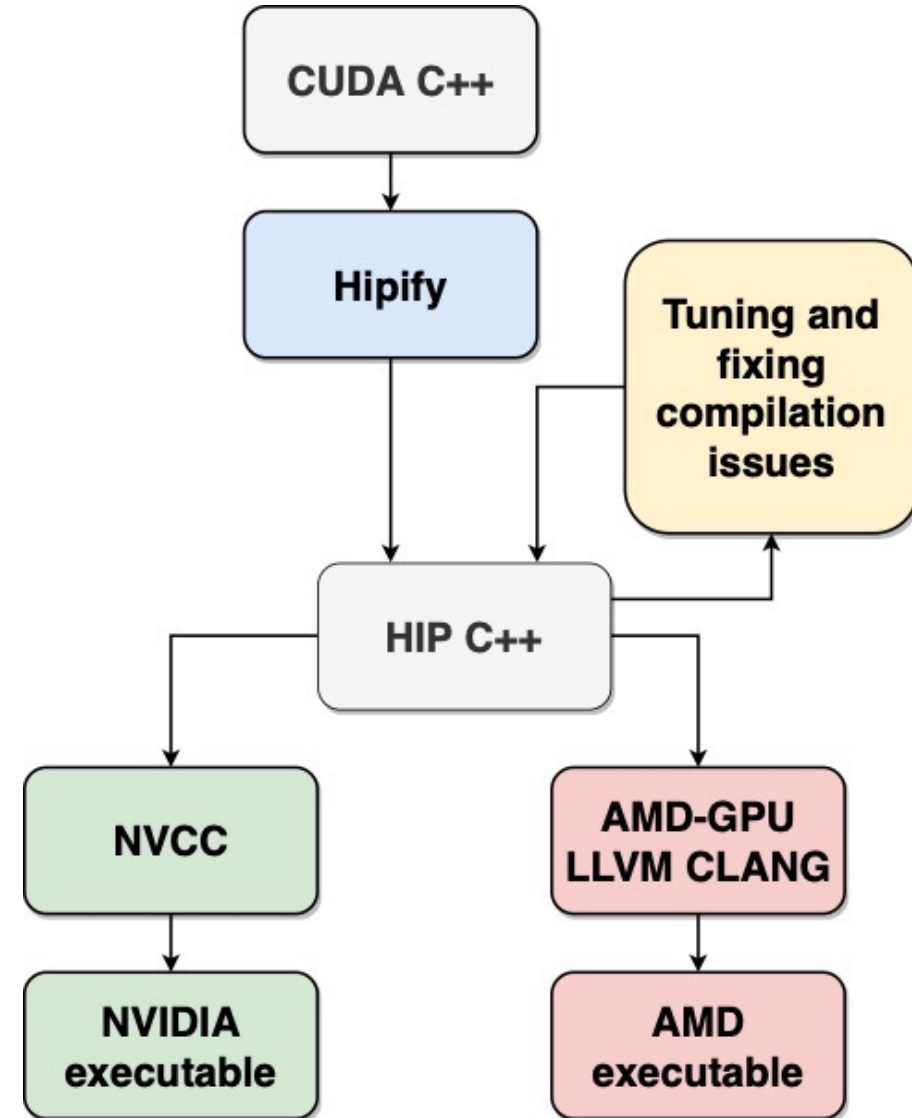
HIP

Portability: code is compiled with NVIDIA or AMD compilers.


Programmability: semantic is similar to CUDA (> 99%).

- First attempt (industry partner from France):
- Porting *all* CUDA code to HIP with Hipify
- Fixed some compilation issues
- Same performance on NVIDIA GPUs
- Poor performance on AMD GPUs

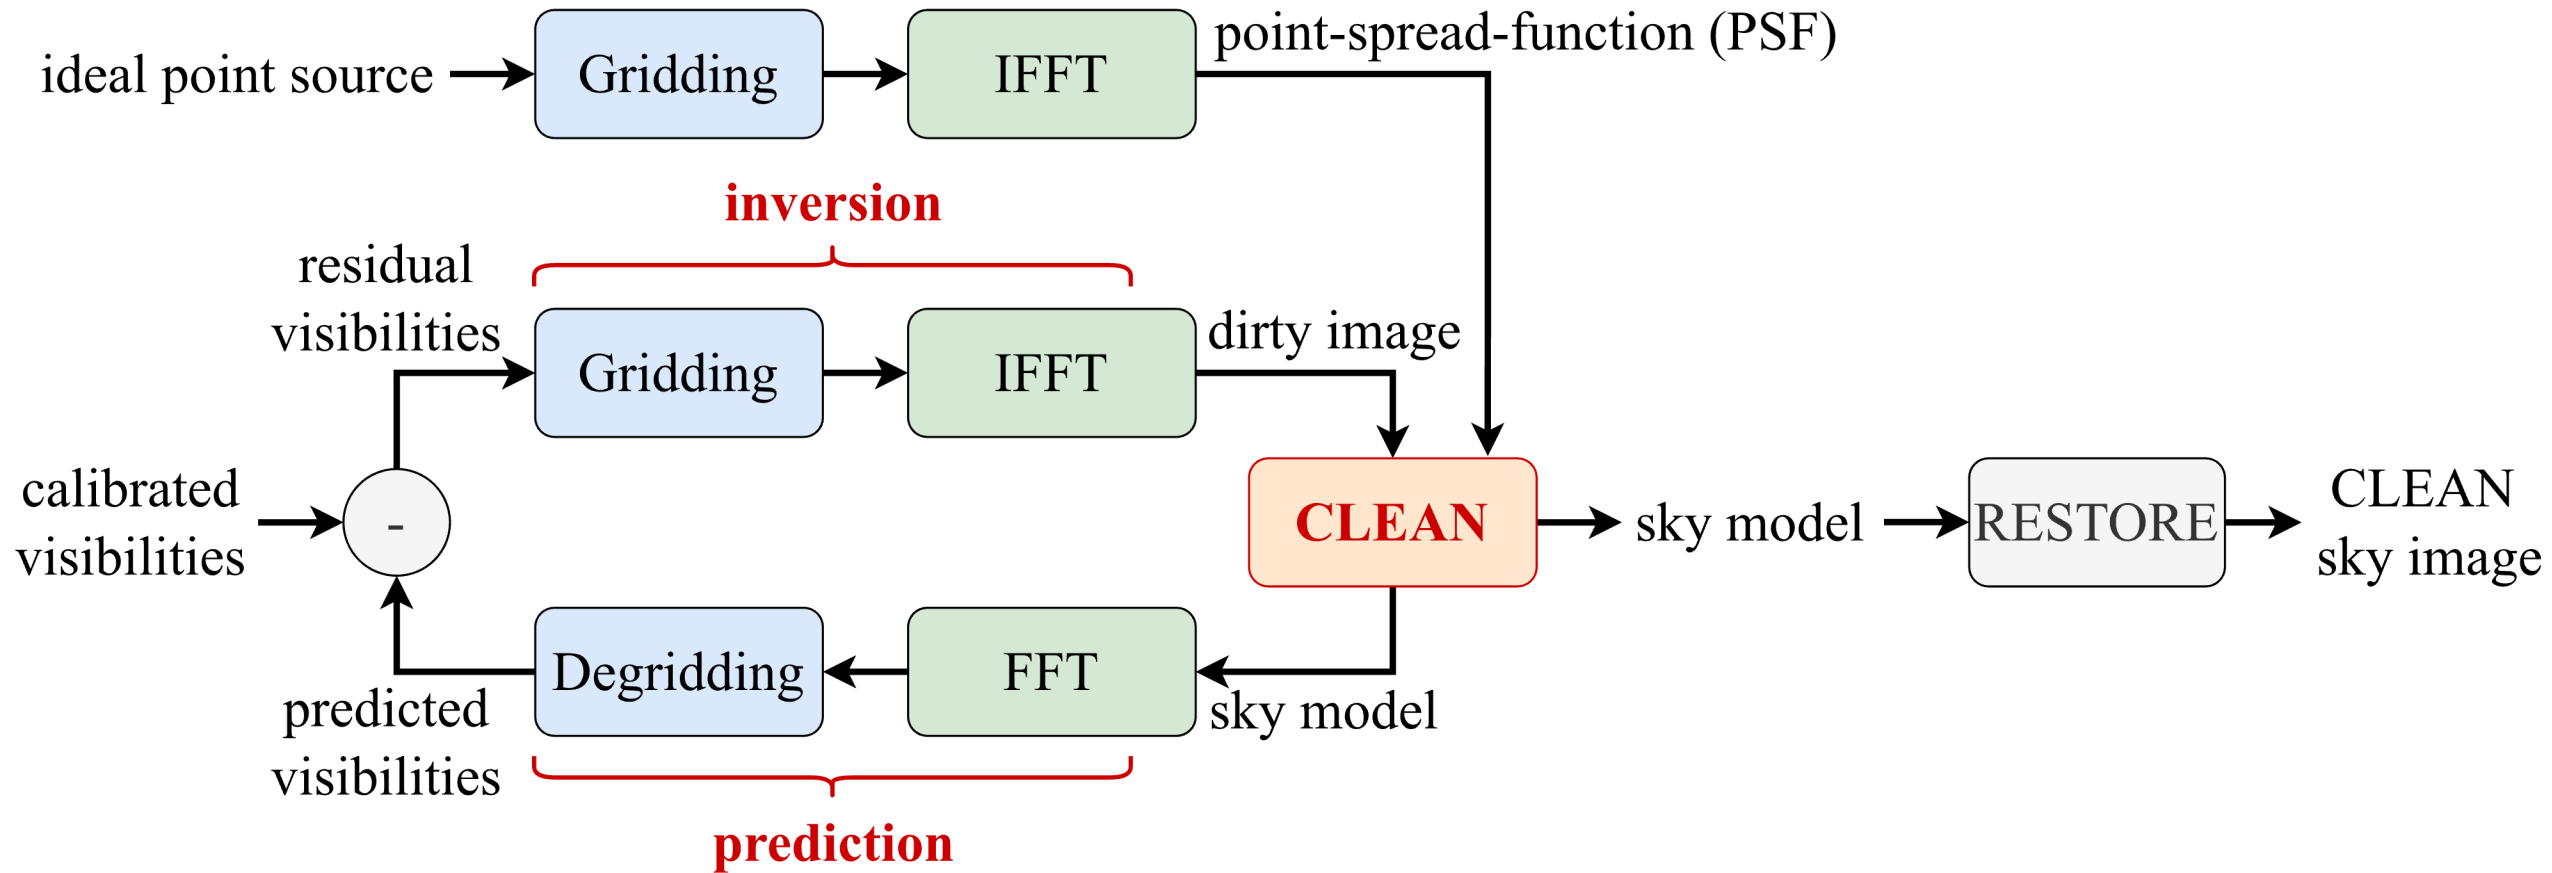
→ **Start again, port one kernel at a time**



Interferometry



Radio-astronomical imaging



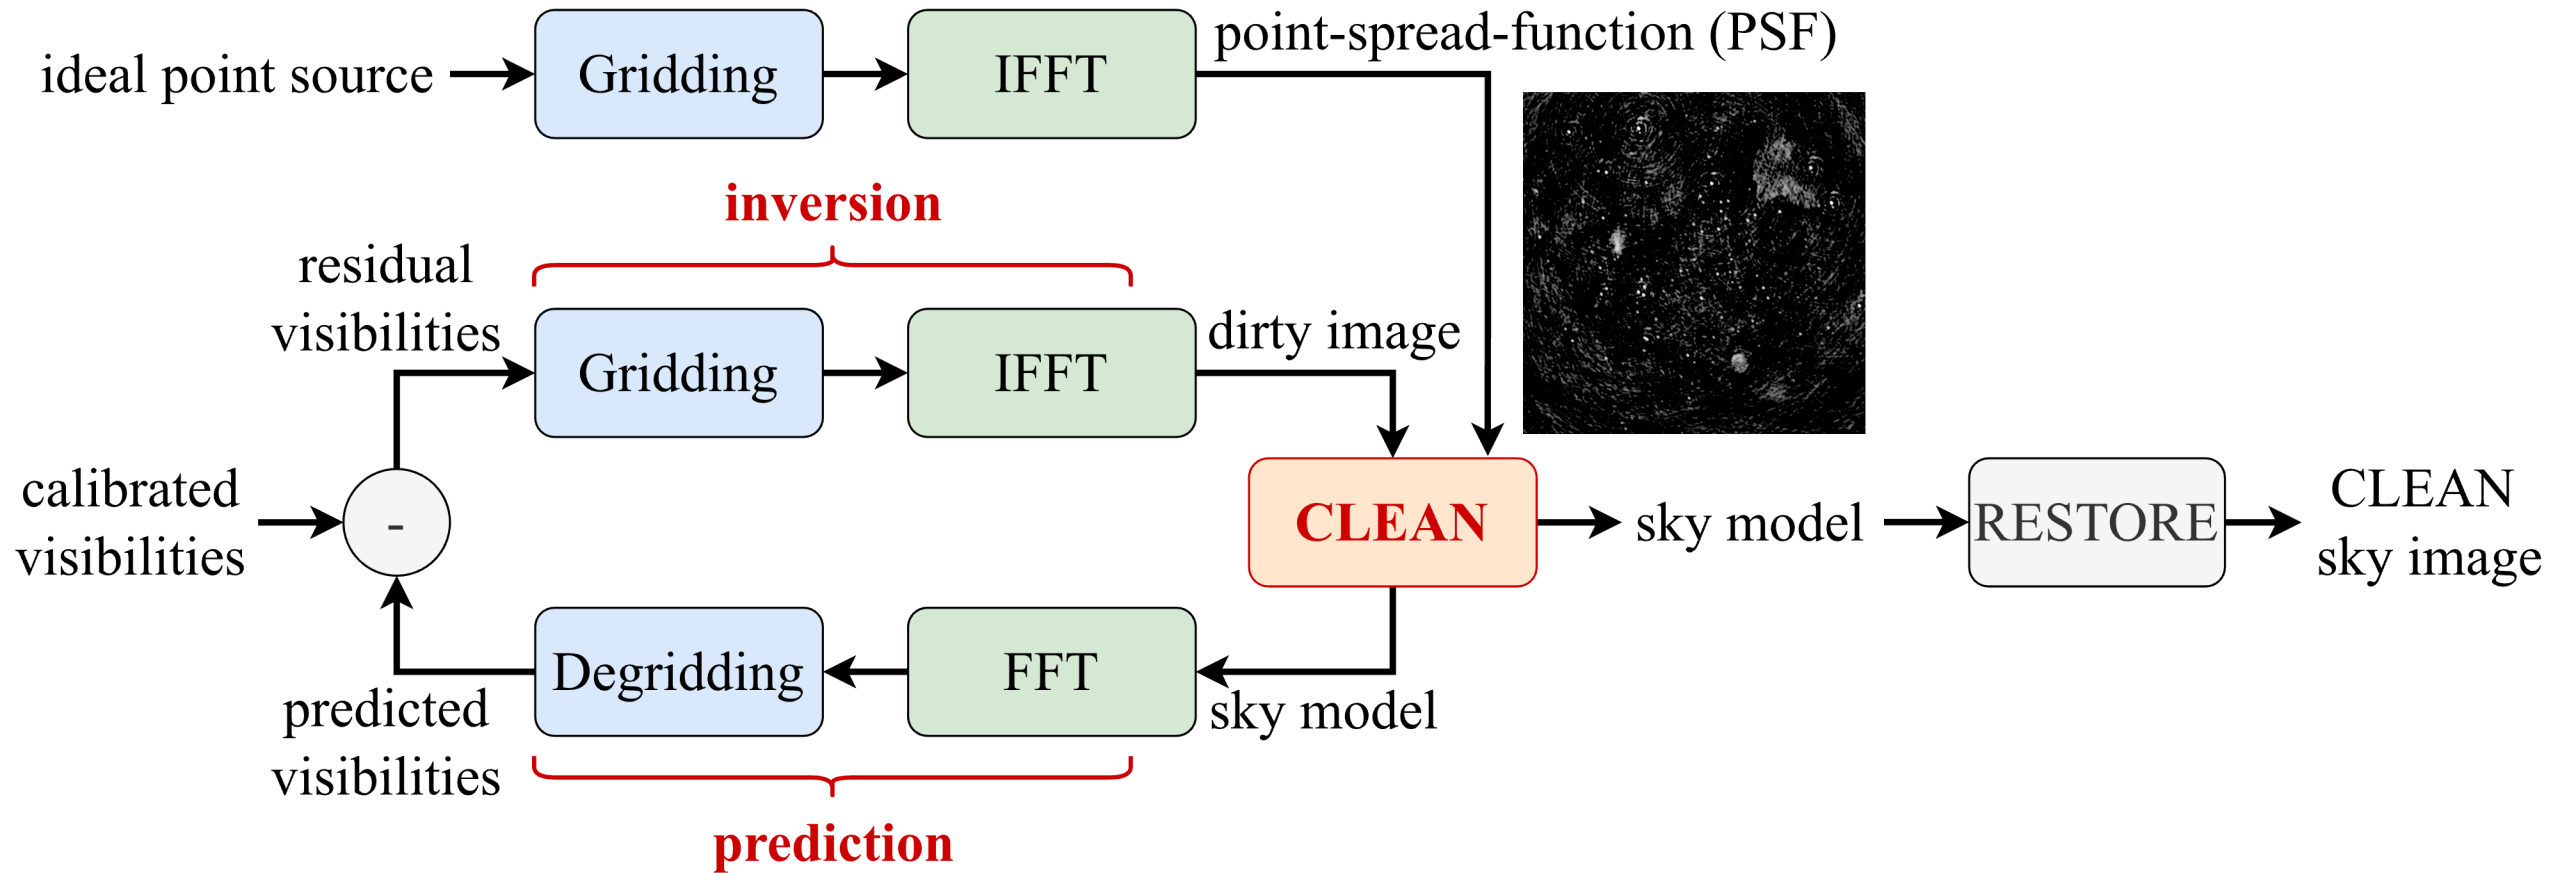
S. Van der Tol et al., "Image Domain Gridding: a fast method for convolutional resampling of visibilities", A&A 2018

A. R. Offringa et al., "An optimized algorithm for multiscale wideband deconvolution of radio astronomical images", MNRAS 2017

S. Corda et al, "Near memory acceleration on high resolution radio astronomy imaging", MECO 2020

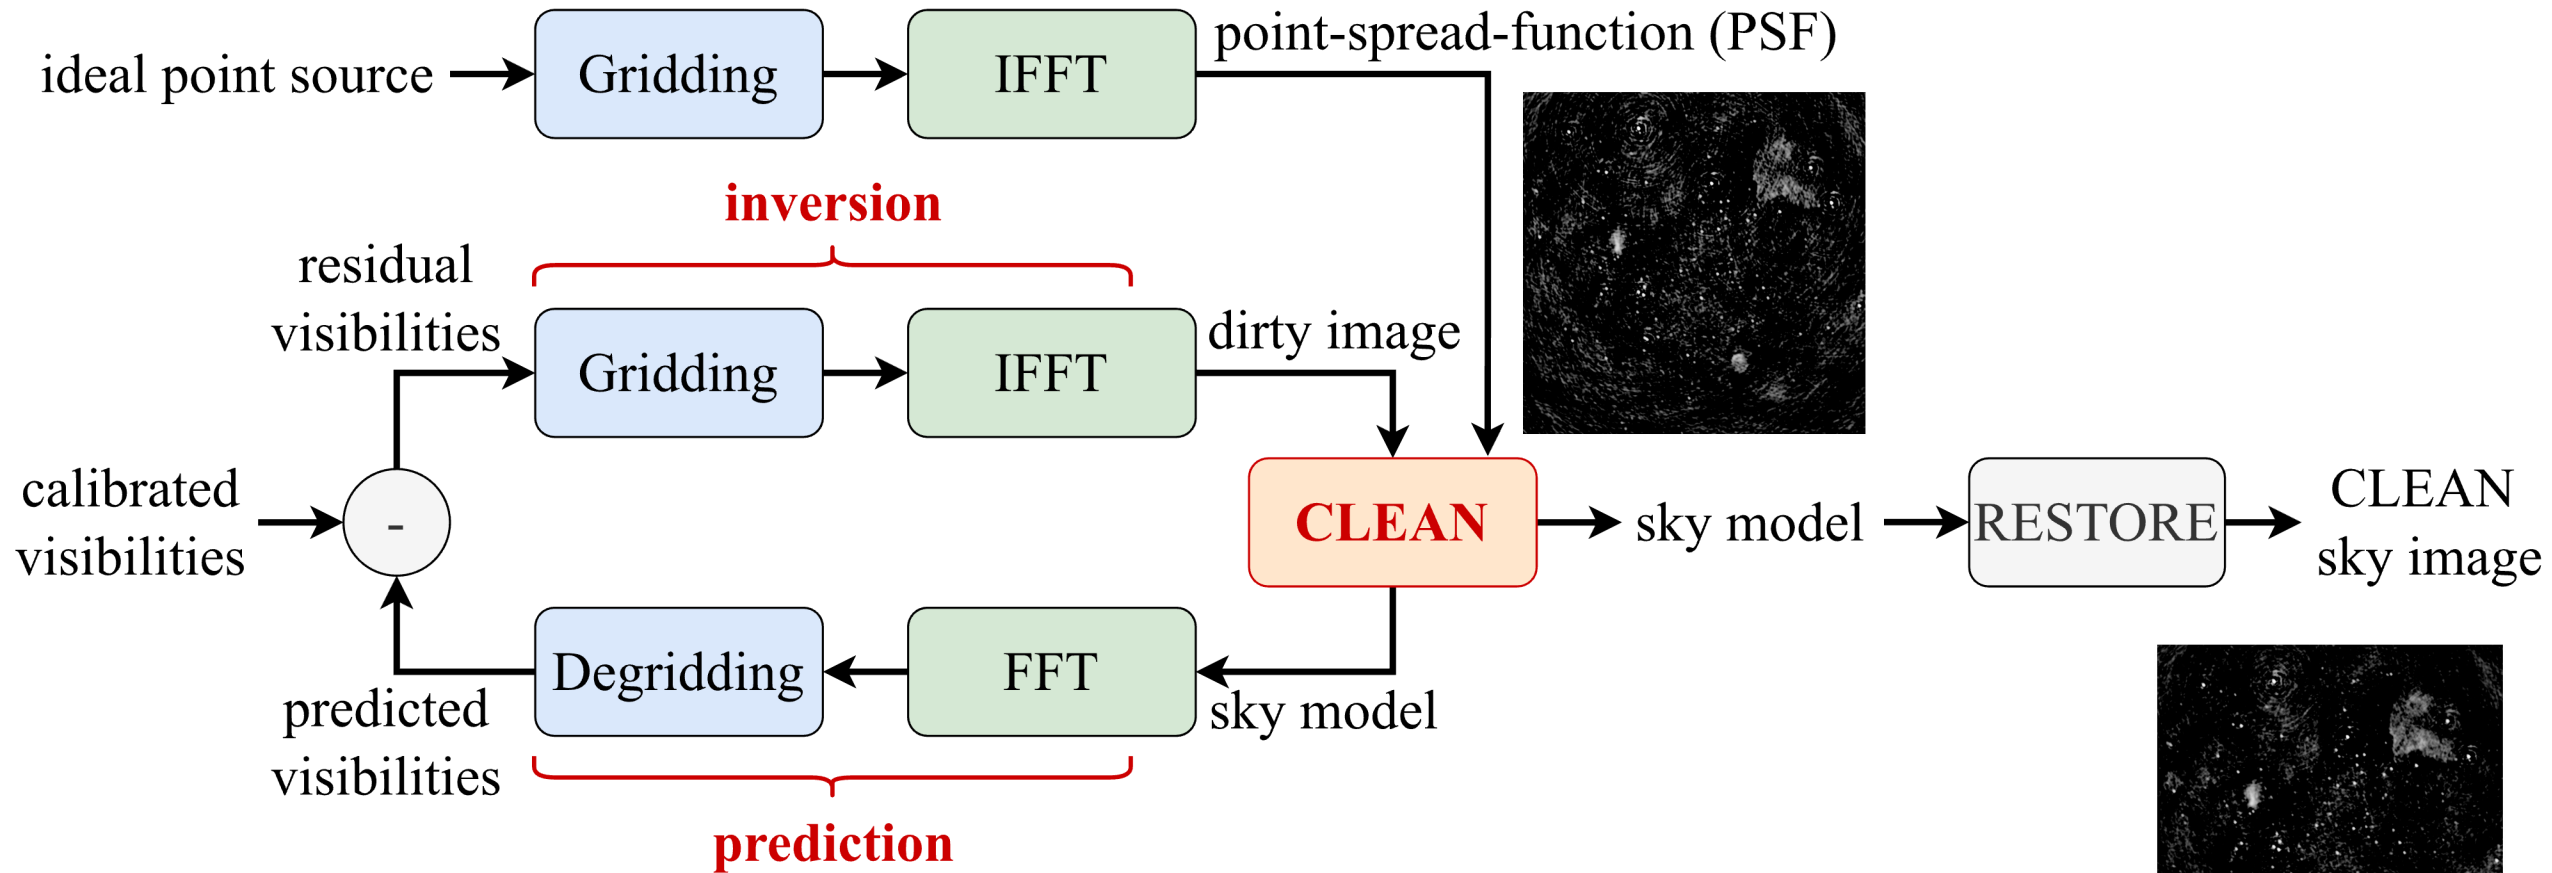
S. Corda et al., "Reduced-precision acceleration of radio-astronomical imaging on reconfigurable hardware", IEEE ACCESS 2022

Radio-astronomical imaging



- S. Van der Tol et al., "Image Domain Gridding: a fast method for convolutional resampling of visibilities", A&A 2018
A. R. Offringa et al., "An optimized algorithm for multiscale wideband deconvolution of radio astronomical images", MNRAS 2017
S. Corda et al, "Near memory acceleration on high resolution radio astronomy imaging", MECO 2020
S. Corda et al., "Reduced-precision acceleration of radio-astronomical imaging on reconfigurable hardware", IEEE ACCESS 2022

Radio-astronomical imaging



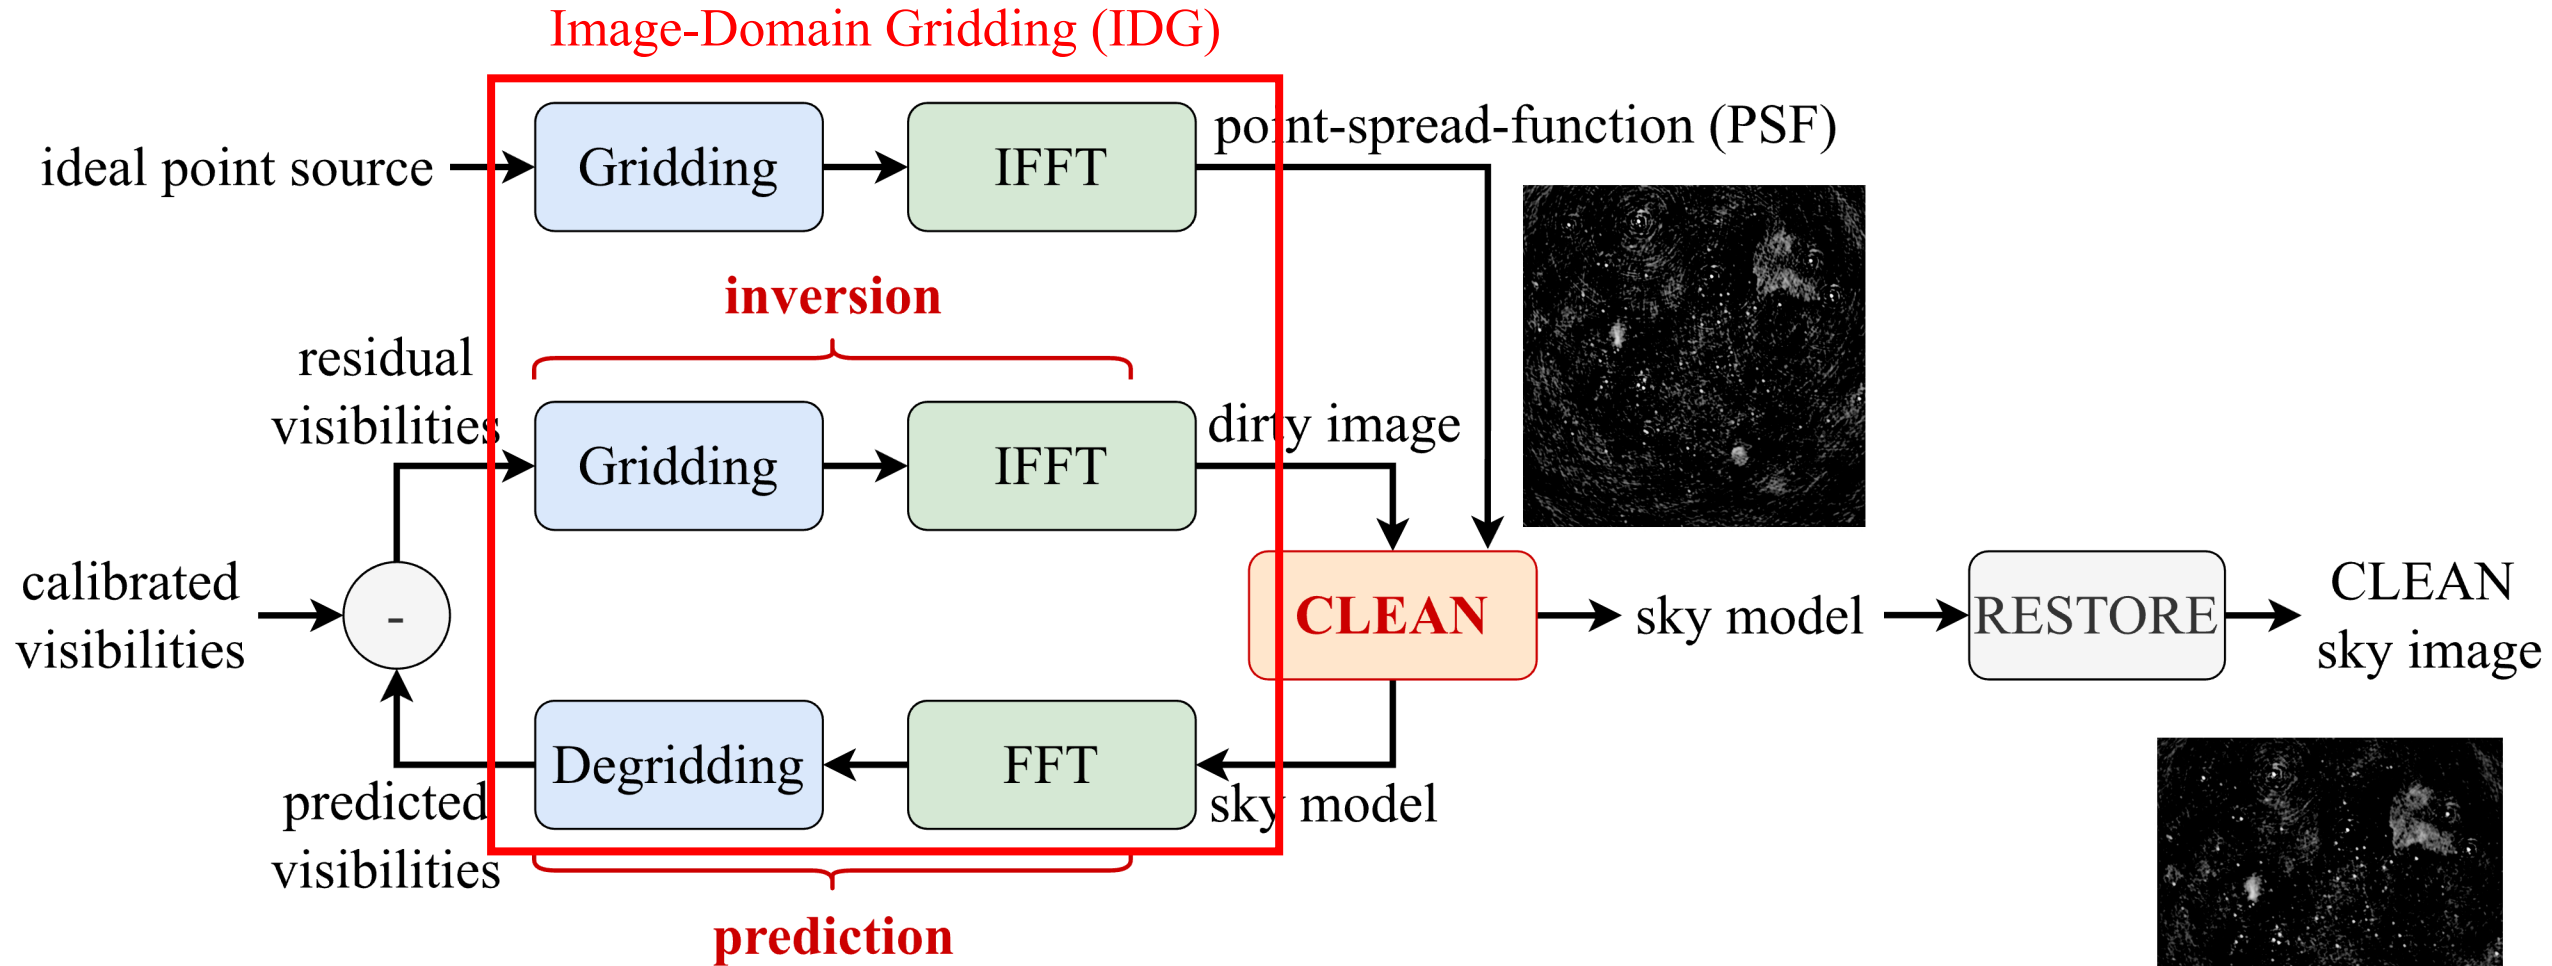
S. Van der Tol et al., "Image Domain Gridding: a fast method for convolutional resampling of visibilities", A&A 2018

A. R. Offringa et al., "An optimized algorithm for multiscale wideband deconvolution of radio astronomical images", MNRAS 2017

S. Corda et al, "Near memory acceleration on high resolution radio astronomy imaging", MECO 2020

S. Corda et al., "Reduced-precision acceleration of radio-astronomical imaging on reconfigurable hardware", IEEE ACCESS 2022

Radio-astronomical imaging



S. Van der Tol et al., "Image Domain Gridding: a fast method for convolutional resampling of visibilities", A&A 2018

A. R. Offringa et al., "An optimized algorithm for multiscale wideband deconvolution of radio astronomical images", MNRAS 2017

S. Corda et al, "Near memory acceleration on high resolution radio astronomy imaging", MECO 2020

S. Corda et al., "Reduced-precision acceleration of radio-astronomical imaging on reconfigurable hardware", IEEE ACCESS 2022

Performance Metrics

The visibilities, which are correlations of station signals, contain information on amplitude and phase of the source.

- **MVis/s**: number of Mega Visibilities per second (throughput)
- **MVis/J**: number of Mega Visibilities per Joule (energy efficiency)

Optimizations

Input: visibilities, wavenumbers, uvw, uvw_offset, lmn

Result: subgrids

```
1 subgrids ← 0;
2 for  $s = 1 \dots S$  do
3   for  $p = 1 \dots NxN$  do
4     complex<float> pixel[pol] ← 0;
5     float lmn [3] ← lmn[p]
6     float phase_offset ← compute_phase_offset(uvw_offsets, lmn)
7     for  $t = 1 \dots T$  do
8       float phase_index ← compute_phase_index(uvw, lmn)
9       for  $c = 1 \dots C$  do
10        float phase ← compute_phase(phase_index, phase_offset, wavenumbers)
11        float phasor [2] ← cosisin(phase)
12        for  $pol$  in polarizations do
13          complex<float> pixel[pol] += visibilities[t][c][pol] * phasor
14  apply_aterm(subgrid);
15  apply_taper(subgrid);
```

Optimizations

v2



NVIDIA A100 whitepaper

Input: visibilities, wavenumbers, uvw, uvw_offset, lmn

Result: subgrids

```
1 subgrids ← 0;
2 for s = 1...S do
3   for p = 1...NxN do
4     complex<float> pixel[pol] ← 0;
5     float lmn [3] ← lmn[p]
6     float phase_offset ← compute_phase_offset(uvw_offsets, lmn)
7     for t = 1...T do
8       float phase_index ← compute_phase_index(uvw, lmn)
9       for c = 1...C do
10        float phase ← compute_phase(phase_index, phase_offset, wavenumbers)
11        float phasor [2] ← cosisin(phase)
12        for pol in polarizations do
13          complex<float> pixel[pol] += visibilities[t][c][pol] * phasor
14 apply_aterm(subgrid);
15 apply_taper(subgrid);
```


Optimizations

v2

v5/v6



NVIDIA A100 whitepaper

```
Input: visibilities, wavenumbers, uvw, uvw_offset, lmn  
Result: subgrids  
1 subgrids  $\leftarrow$  0;  
2 for  $s = 1 \dots S$  do  
3   for  $p = 1 \dots NxN$  do  
4     complex<float> pixel[pol]  $\leftarrow$  0;  
5     float lmn [3]  $\leftarrow$  lmn[p]  
6     float phase_offset  $\leftarrow$  compute_phase_offset(uvw_offsets, lmn)  
7     for  $t = 1 \dots T$  do  
8       float phase_index  $\leftarrow$  compute_phase_index(uvw, lmn)  
9       for  $c = 1 \dots C$  do  
10        float phase  $\leftarrow$  compute_phase(phase_index, phase_offset, wavenumbers)  
11        float phasor [2]  $\leftarrow$  cosisin(phase)  
12        for  $pol$  in polarizations do  
13          complex<float> pixel[pol] += visibilities[t][c][pol] * phasor  
14    apply_aterm(subgrid);  
15    apply_taper(subgrid);
```

Optimizations

v2

v5/v6



NVIDIA A100 whitepaper

```

Input: visibilities, wavenumbers, uvw, uvw_offset, lmn
Result: subgrids
1 subgrids ← 0;
2 for s = 1...S do
3   for p = 1...NxN do
4     complex<float> pixel[pol] ← 0;
5     float lmn [3] ← lmn[p]
6     float phase_offset ← compute_phase_offset(uvw_offsets, lmn)
7     for t = 1...T do
8       float phase_index ← compute_phase_index(uvw, lmn)
9       for c = 1...C do
10        float phase ← compute_phase(phase_index, phase_offset, wavenumbers)
11        float phasor [2] ← cosisin(phase)
12        for pol in polarizations do
13          complex<float> pixel[pol] += visibilities[t][c][pol] * phasor
14      apply_aterm(subgrid);
15      apply_taper(subgrid);
    
```

Original

v9

Modified to use Packed FMA32

```

float vxi = 0.0f, vyi = 0.0f, vzi = 0.0f;
for (int j = hipThreadIdx_x; j < count1; j += hipBlockDim_x) {
    float dx = xx1[j] - xxi;
    float dy = yy1[j] - yyi;
    float dz = zz1[j] - zzi;
    float dist2 = dx*dx + dy*dy + dz*dz;
    if (dist2 < fsrrmax2) {
        float rtemp = (dist2 + rsm2)*(dist2 + rsm2)*(dist2 + rsm2);
        float f_over_r = massi*mass1[j]*(1.0f/sqrt(rtemp) - (ma0 +
        dist2*(ma1 + dist2*(ma2 + dist2*(ma3 + dist2*(ma4 + dist2*ma5))))));

        vxi += fcoeff*f_over_r*dx;
        vyi += fcoeff*f_over_r*dy;
        vzi += fcoeff*f_over_r*dz;
    }
}
    
```

```


float2 vxi = 0.0f, vyi = 0.0f, vzi = 0.0f;
for (int j = hipThreadIdx_x; j < count1; j += 2*hipBlockDim_x) {
    float2 dx = {xx1[j] - xxi, xx1[j+ hipBlockDim_x] - xxi};
    float2 dy = {yy1[j] - yyi, yy1[j+ hipBlockDim_x] - yyi};
    float2 dz = {zz1[j] - zzi, zz1[j+ hipBlockDim_x] - zzi};
    float2 dist2 = dx*dx + dy*dy + dz*dz;
    if (dist2 < fsrrmax2) {
        float2 rtemp = (dist2 + rsm2)*(dist2 + rsm2)*(dist2 + rsm2);
        float2 f_over_r = massi*mass1[j]*(1.0f/sqrt(rtemp) - (ma0 +
        dist2*(ma1 + dist2*(ma2 + dist2*(ma3 + dist2*(ma4 + dist2*ma5))))));

        vxi += fcoeff*f_over_r*dx;
        vyi += fcoeff*f_over_r*dy;
        vzi += fcoeff*f_over_r*dz;
    }
}
    
```

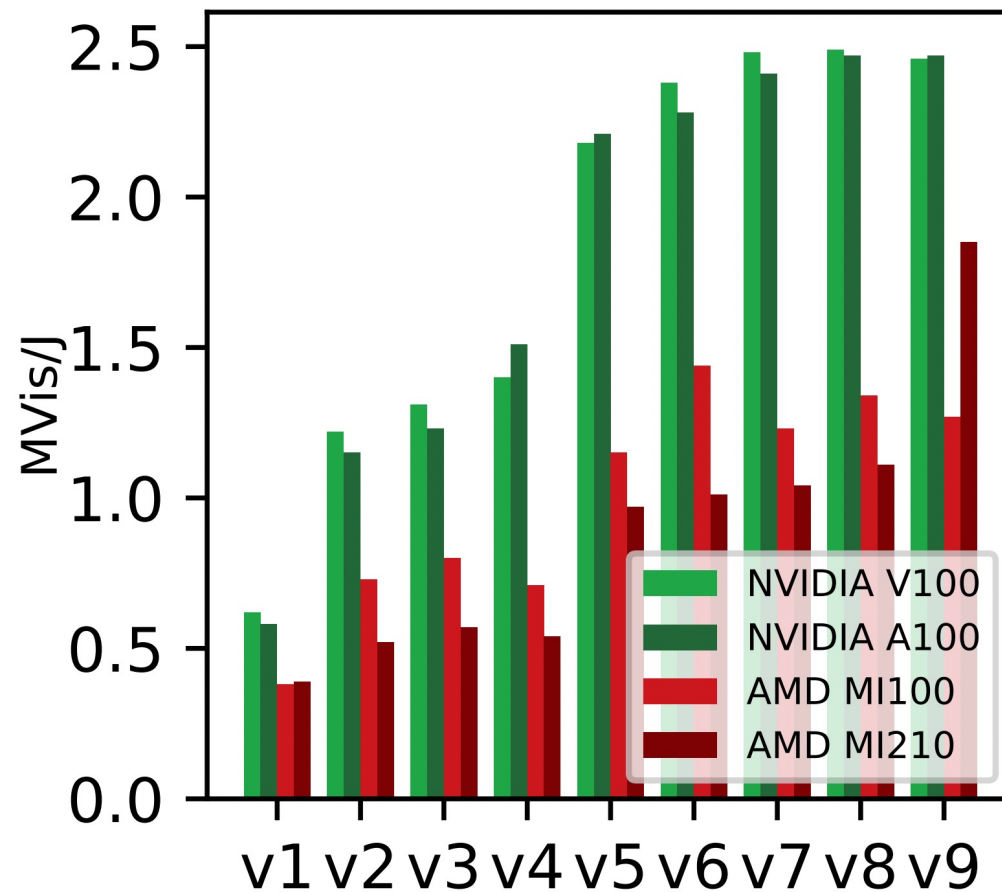
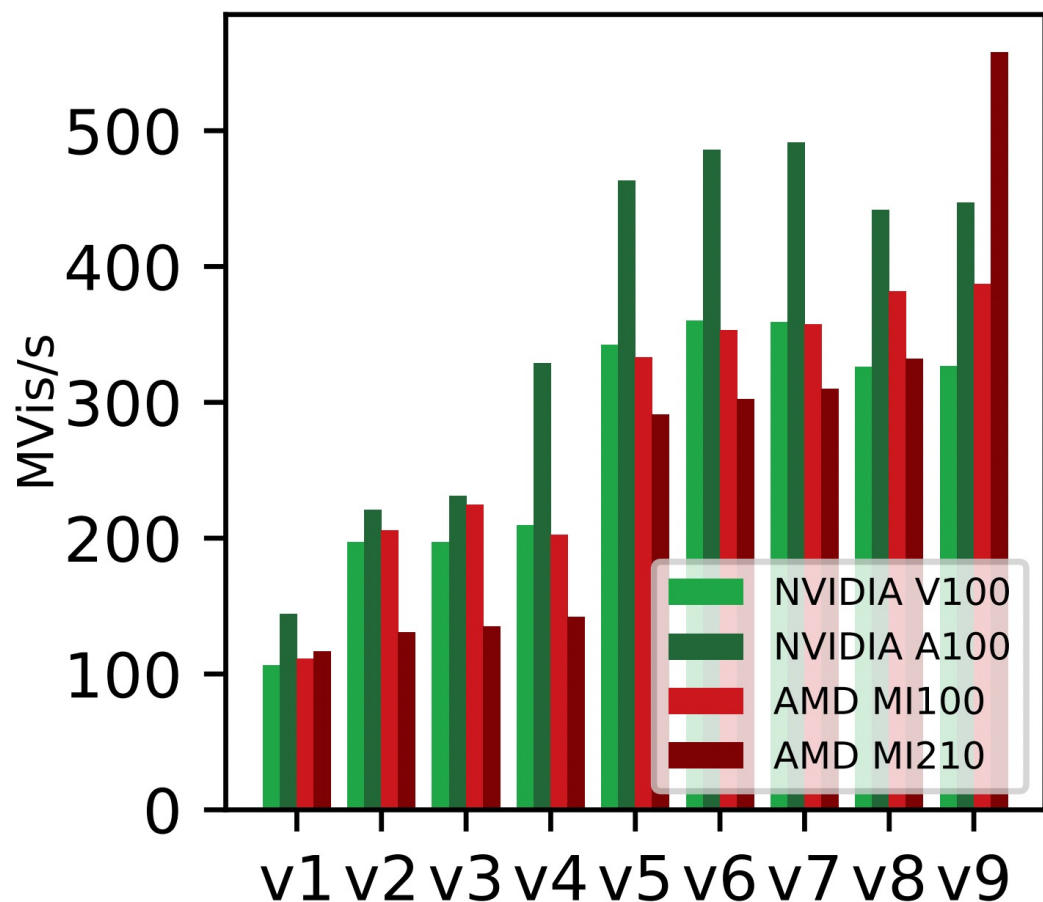
AMD CNDA2 whitepaper

V100: CUDA vs HIP

- Marginal differences
- Certain optimization are better implemented by HIP and others by NVIDIA compiler




Performance (gridder)




- CDNA2 shows good potential by exploiting FP32 packed instructions (v9)
- Application code needs additional tuning to achieve better performance

Large 2D FFTs




■ “Cold run”
issue

Large 2D FFTs



■ “Cold run”
issue



■ GPUs
outperform
CPUs

Conclusions

- HIP provides equivalent performance for the IDG gridder/degridder on NVIDIA GPUs
- HIP provides "acceptable" performance portability to AMD GPUs
 1. Easier than maintaining both a CUDA and OpenCL codebase
 2. Additional tuning is needed to get closer to peak performance
- Recent AMD GPUs can achieve competitive performance if the application can be re-shaped to support single-precision packed instructions
- We step on several issues running HIPFFT (performance on AMD and compatibility with NVIDIA).

Conclusions

- HIP provides equivalent performance for the IDG gridder/degridder on NVIDIA GPUs
- HIP provides "acceptable" performance portability to AMD GPUs
 1. Easier than maintaining both a CUDA and OpenCL codebase
 2. Additional tuning is needed to get closer to peak performance
- Recent AMD GPUs can achieve competitive performance if the application can be re-shaped to support single-precision packed instructions
- We step on several issues running HIPFFT (performance on AMD and compatibility with NVIDIA).
- Next Steps:
 - Gridder/Degridder performance difference analysis and optimization
 - Discuss with AMD HIPFFT issues
 - Port CUDA production code with HIP
 - Co-design evaluation on several AMD/NVIDIA GPU generations

Radio-astronomical imaging on GPUs and FPGAs

Thanks!

Stefano Corda

stefano.corda@epfl.ch

04-10-2022

Swiss SKA Days 2022, Lugano

