


NRAO Algorithms R&D Group

DPUC, March 19th 2025




S. Bhatnagar

Algorithms R&D Group,
National Radio Astronomy Observatory, Socorro, NM, USA



Drivers


- Cognizant of inherent complexity



- Evidence driven R&D
 - Adapt to algorithms and technology evolution: Scalable algorithm and software architectures
 - Scaling on large heterogeneous platforms : Simplify the tool-chain
 - Address the higher compute load : Integrated algorithm development
 - Need to involve other researchers in the community as content producers
- Collaborations
 - RA projects/groups interested in utilizing *full* telescope capabilities
 - With HTC/HPC/Industry: Computing at the necessary scale

Current work

- Problem: Address issues due to next-gen size of computing



!!

- ngVLA: 10s PFLOP/sec ($T_{\text{rec}} \sim 1$ day) \leftarrow ~~$O(\text{million})$ CPU cores~~ / few x $O(1000\text{s})$ GPUs
- WSU: $O(100)$ TFLOP/sec \leftarrow Can be done today!

https://library.nrao.edu/public/memos/ngvla/NGVLAC_04.pdf

- Solution: Scalable architecture
 - Flexibility required to benefit from new algorithms, technology evolution, collaborations with others groups,...
- Solution: New Algorithms:
 - Can new algorithms help with compute scaling?
 - Address predominantly “high” frequency issues: Water vapor, Antenna pointing errors (ngVLA)

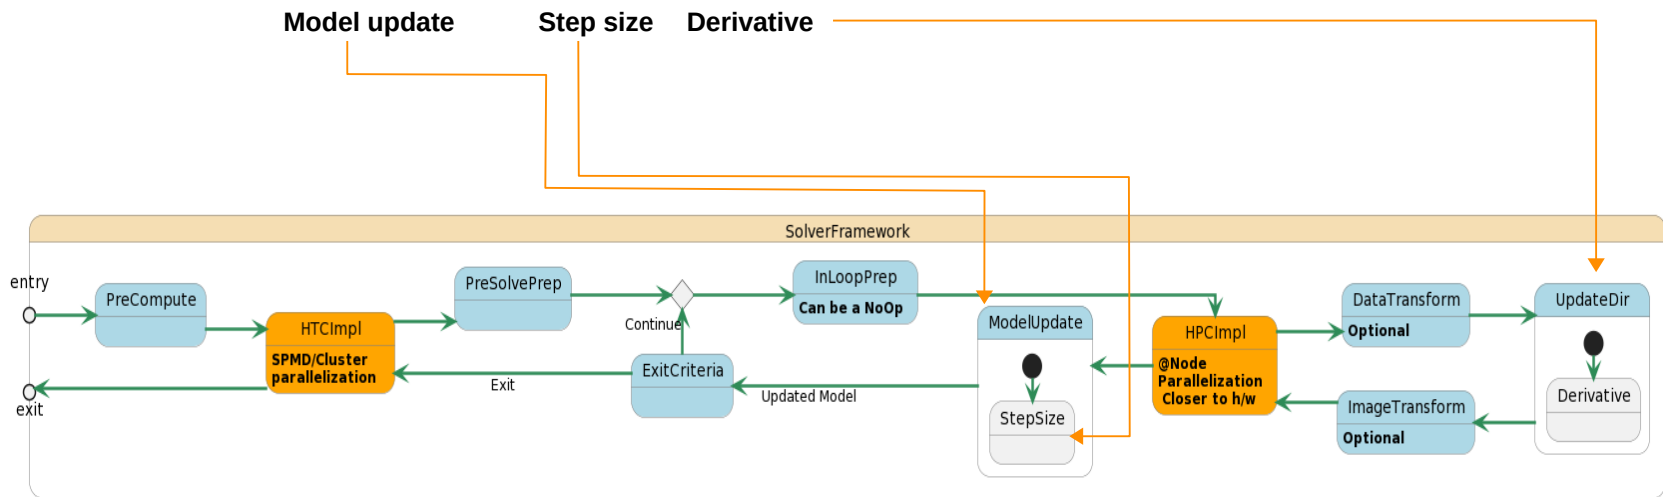
Scalable Algorithm Architecture

- Based on mathematical, signal processing, domain (rather than _____) considerations
 - Evolving computing needs, algorithms, h/w & s/w
 - Separable components that scale individually and together

$$V^{obs} = \mathbf{G}^M \mathbf{S} \mathbf{F} \mathbf{B}^M \mathbf{I}^M + noise$$

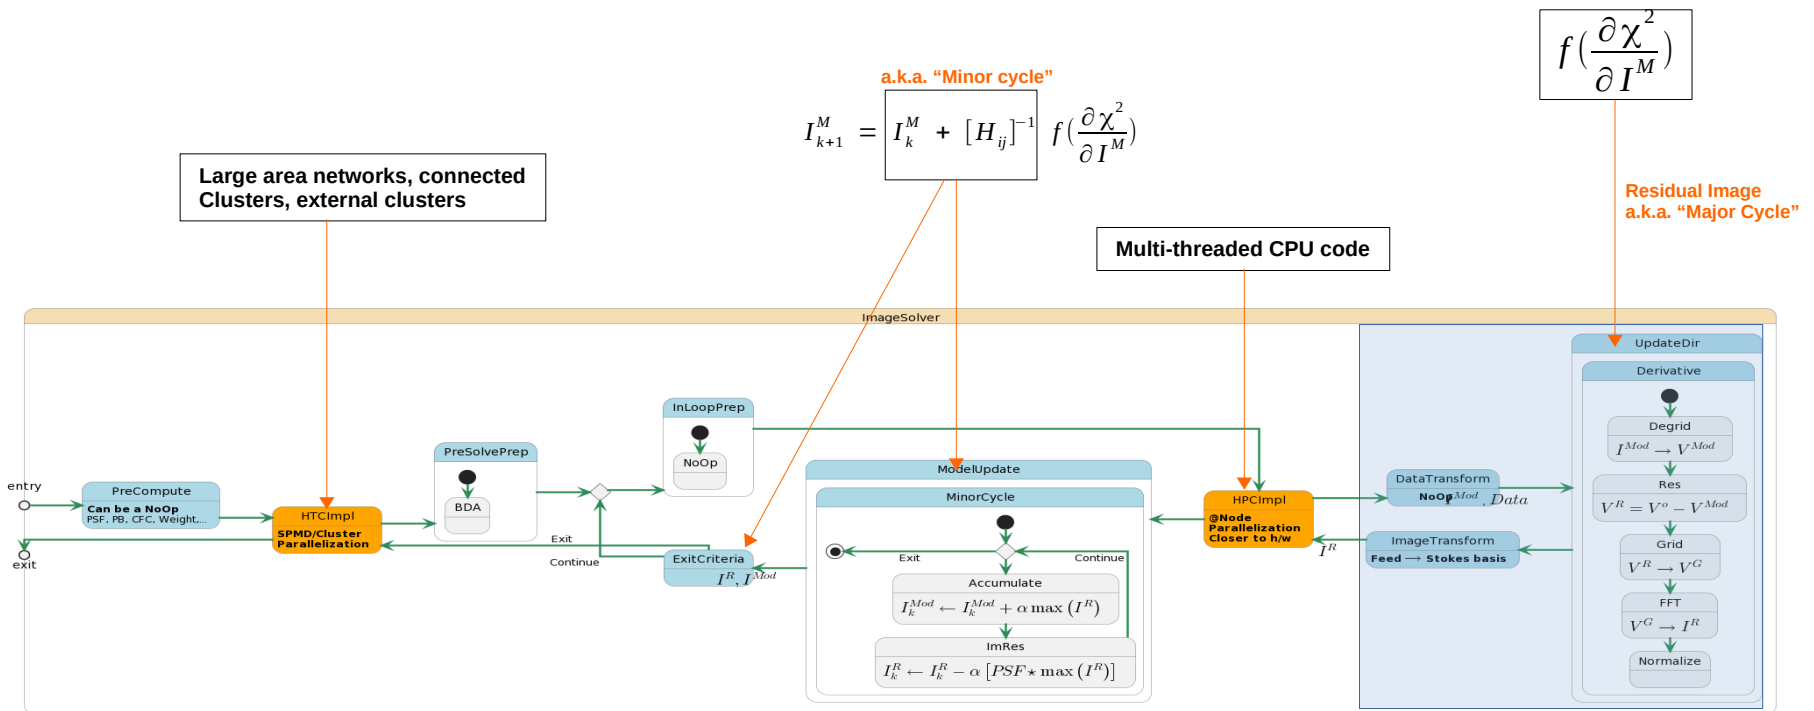
$$\chi^2 = \sum_i |Data_i - Model_i(P)|^2$$

$$P_i^{k+1} = P_i^k + [H_{ij}]^{-1} f\left(\frac{\partial \chi^2}{\partial P_i^k}\right) ; \quad [H_{ij}] = \frac{\partial^2 \chi^2}{\partial P_i^k \partial P_j^k}$$



Algorithm Architecture: Imaging

- Component specialization delivers various calibration and imaging algorithms



LibRA: User view

- Algorithms (current & future) can be expressed in a framework with a collection of simpler, fundamental components
- Components independently configurable via multiple interfaces

<https://github.com/ARDG-NRAO/LibRA>

LibRA: User view

- Algorithms (current & future) can be expressed in a framework with a collection of simpler, fundamental components
- Components independently configurable via multiple interfaces

Command-line

```
roadrunner help=noprompt vis="..." imagename="..." isize="12*1024"
```

Python

```
>>> import roadrunner2py  
>>> roadrunner2py.Roadrunner(vis="...", imagename="...", isize="12*1024"
```

Library

```
#include <roadrunner.h>  
  
int main(...)  
{  
:  
:  
  RRReturnType rrType;  
  rrType=Roadrunner(visName, imageName,  
:  
:  
}
```

<https://github.com/ARDG-NRAO/LibRA>

LibRA: User view

- Algorithms (current & future) can be expressed in a framework with a collection of simpler, fundamental components
- Components independently configurable via multiple interfaces

Command-line

```
roadrunner help=noprompt vis="..." imagename="..." isize="12*1024"
```

Python

```
>>> import roadrunner2py  
>>> roadrunner2py.Roadrunner(vis="...", imagename="...", isize="12*1024")
```

Library

```
#include <roadrunner.h>  
  
int main(...)  
{  
    :  
    :  
    RRReturnType rrType;  
    rrType=Roadrunner(visName, imageName,  
    :  
    :  
}
```

- Deployed on a wide range of computing resources


Laptop



Desktop



Cluster of GPUs/CPU



Super-computer class

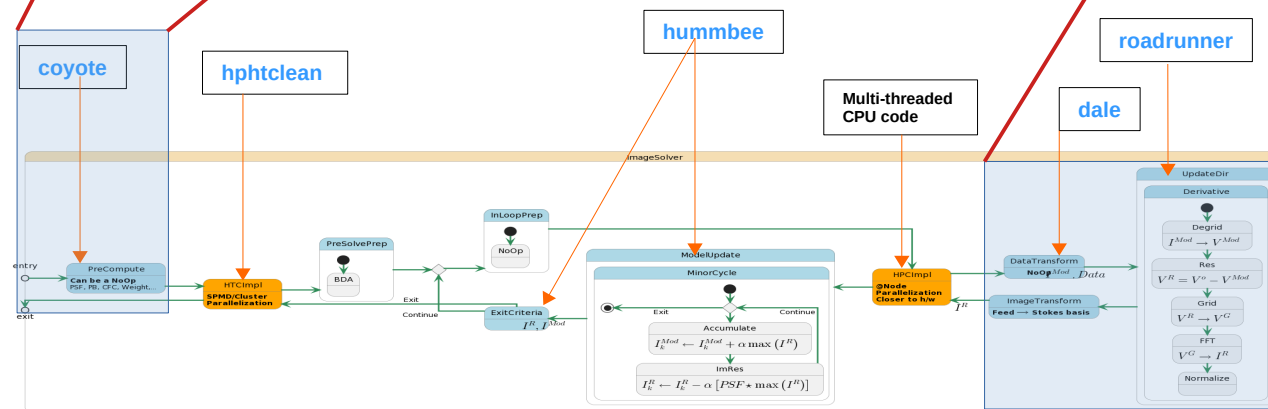


<https://github.com/ARDG-NRAO/LibRA>

Unlock access to *existing* resources


- Significant computing resources *already* available (US community).
 - PATH facility: The Center for High Throughput Computing (CHTC, UW-M)
 - National Research Platform (NRP) + The San Diego Super Computer Center (UCSD + SDSC)

Domain specialization



Unlock access to *existing* resources


- Deployed on a cluster of GPUs (~100) on the PATH facility in collaboration with
 - PATH facility: The Center for High Throughput Computing (CHTC, UW-M)
 - National Research Platform (NRP) + The San Diego Super Computer Center (UCSD + SDSC)



<https://science.nrao.edu/enews/17.3/index.shtml#deepimaging>

Simplify access

- Deployed on a cluster of GPUs (~100) on the PATH facility in collaboration with
 - PATH facility: The Center for High Throughput Computing (CHTC, UW-M)
 - National Research Platform (NRP) + The San Diego Super Computer Center (UCSD + SDSC)



<https://science.nrao.edu/enews/17.3/index.shtml#deepimaging>

Design verification

- Deployed on a cluster of GPUs (~100) on the PATH facility of the Open Science Grid (OSG) in collaboration with <https://science.nrao.edu/enews/17.3/index.shtml#deepimaging>
 - PATH facility: The Center for High Throughput Computing (CHTC, UW-M)
 - National Research Platform (NRP) + The San Diego Super Computer Center (SDSC)

Domain specialization

- Throughput: ~1 TB/hr (~2 TB/hr now?)
- 10 iterations in ~24 hr
 - Earlier attempts using CPU cores:
~14 days per cycle
- This is still a small fraction of the **required peak** throughput (?!)

Among the deepest radio image

- Deployed on a cluster of GPUs (~100) on the PATH and NRP facility




Image made using OSG/PATH and OSDF

Data: Courtesy The CHILES Team

Paper in prep.