

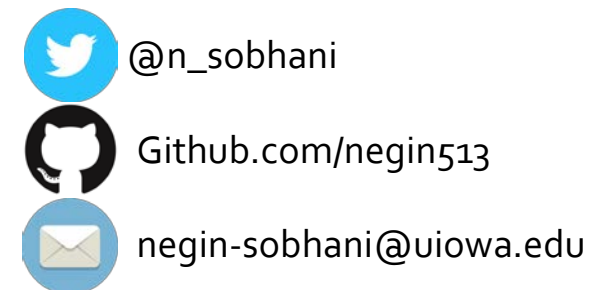


# Performance Analysis and Optimization of the Weather Research and Forecasting Model (WRF) Advection Schemes

Negin Sobhani <sup>1,2</sup>, Davide Del Vento<sup>2</sup>, and Dave Gill<sup>2</sup>

<sup>1</sup>University of Iowa

<sup>2</sup>National Center for atmospheric Research(NCAR)



# WRF is a numerical weather prediction system designed for both atmospheric research and operational forecasting.

Community model with **large user base**:

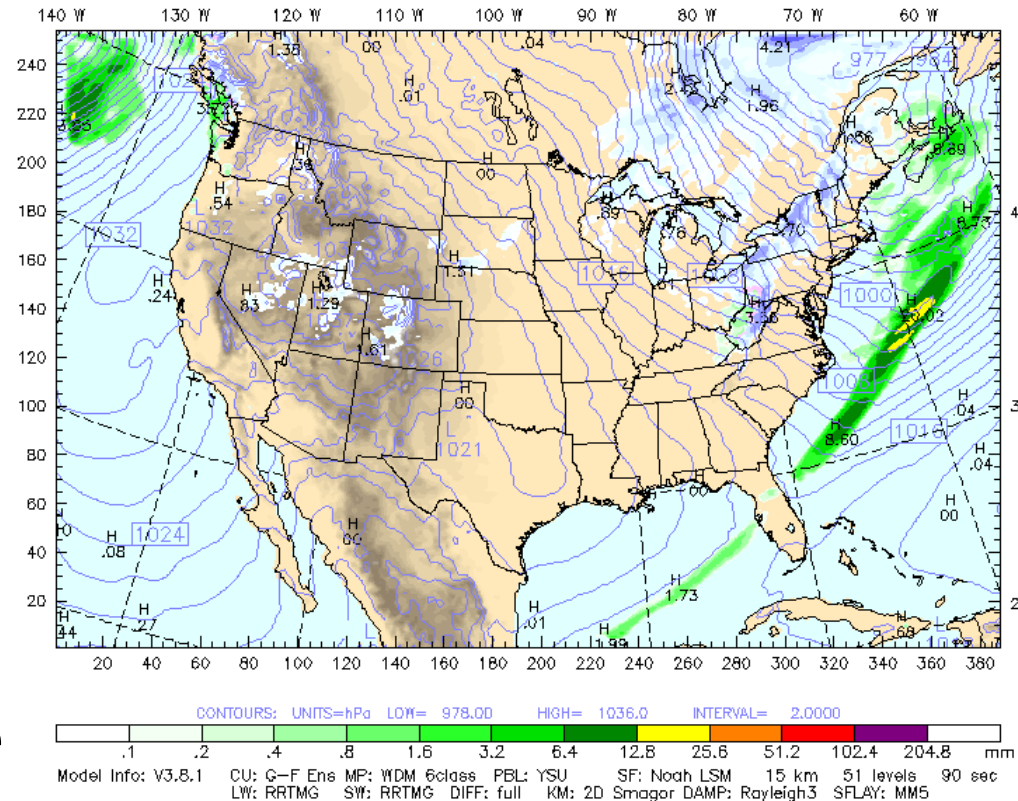
- More than 30,000 users in 150 countries

## Research Goals

**Goal 1.** WRF scalability and comparison between MPI and hybrid parallelism

**Goal 2.** Identify hotspots and potential areas for improvement in WRF

**Goal 3.** Optimize the code to improve the performance of the hotspots



15 km X 15 km

~2 million  
gridpoints

Higher  
resolutions???

# Scalability Assessment (MPI Only)

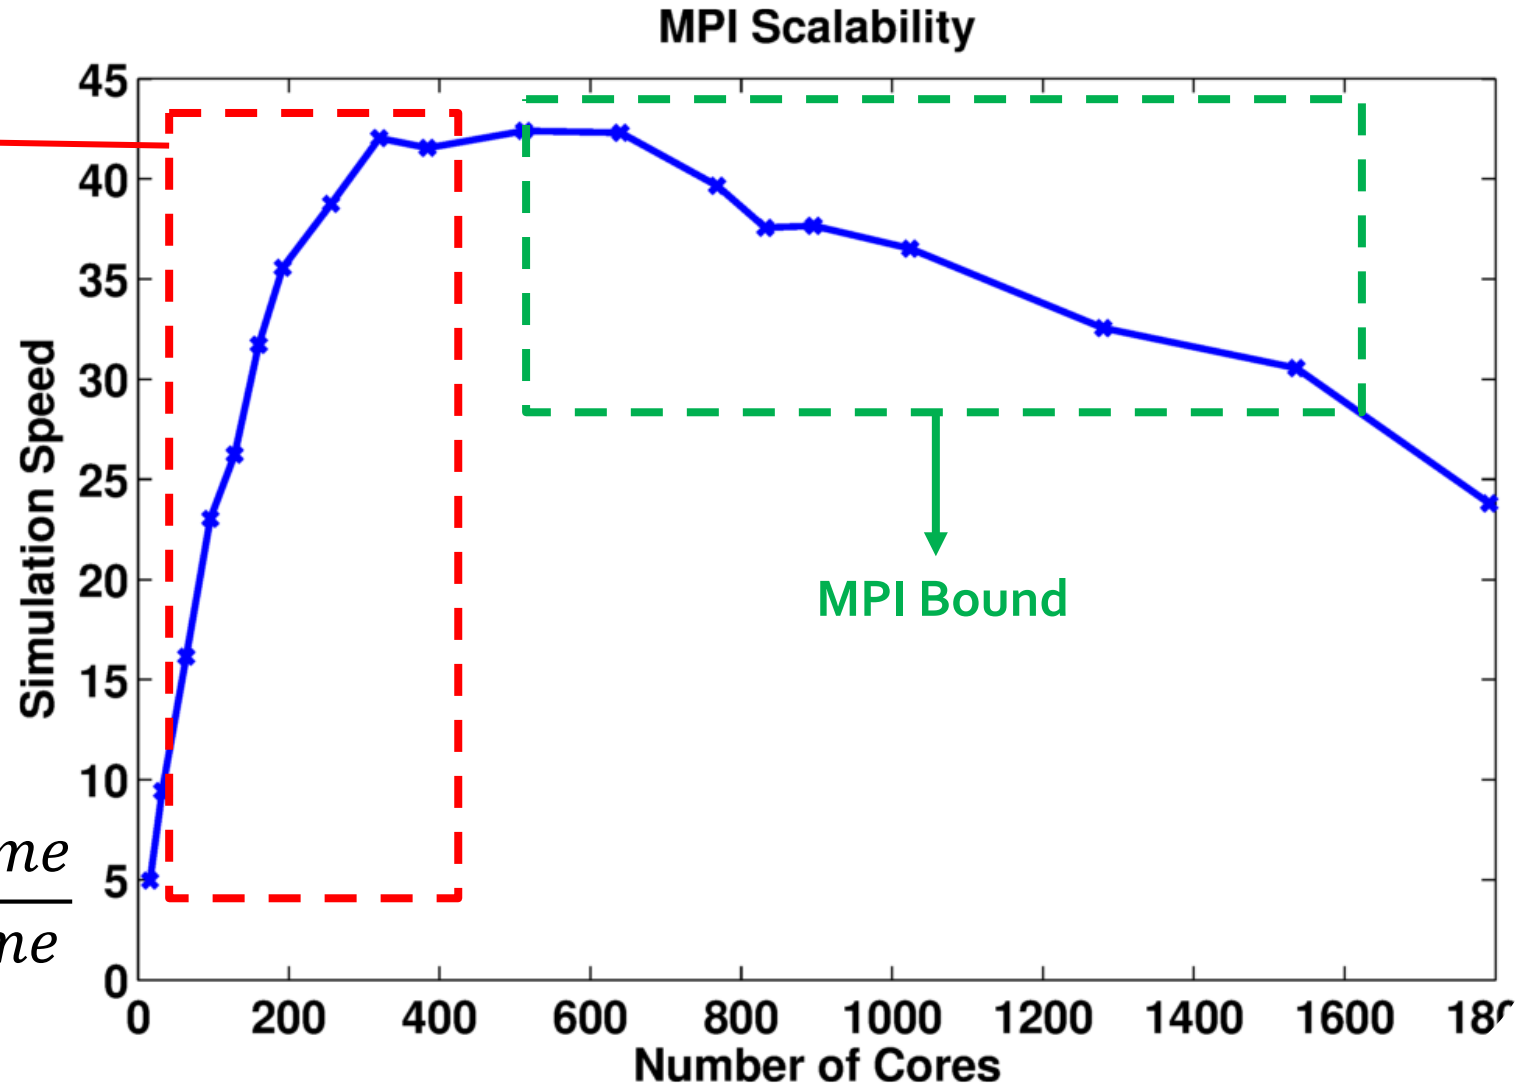
3

500X500 horizontal grids

Compute Bound

Simulation Speed is  
duration of simulation  
per wall clock time

$$\text{Simulation Speed} = \frac{\text{Simulation Time}}{\text{WallClock Time}}$$

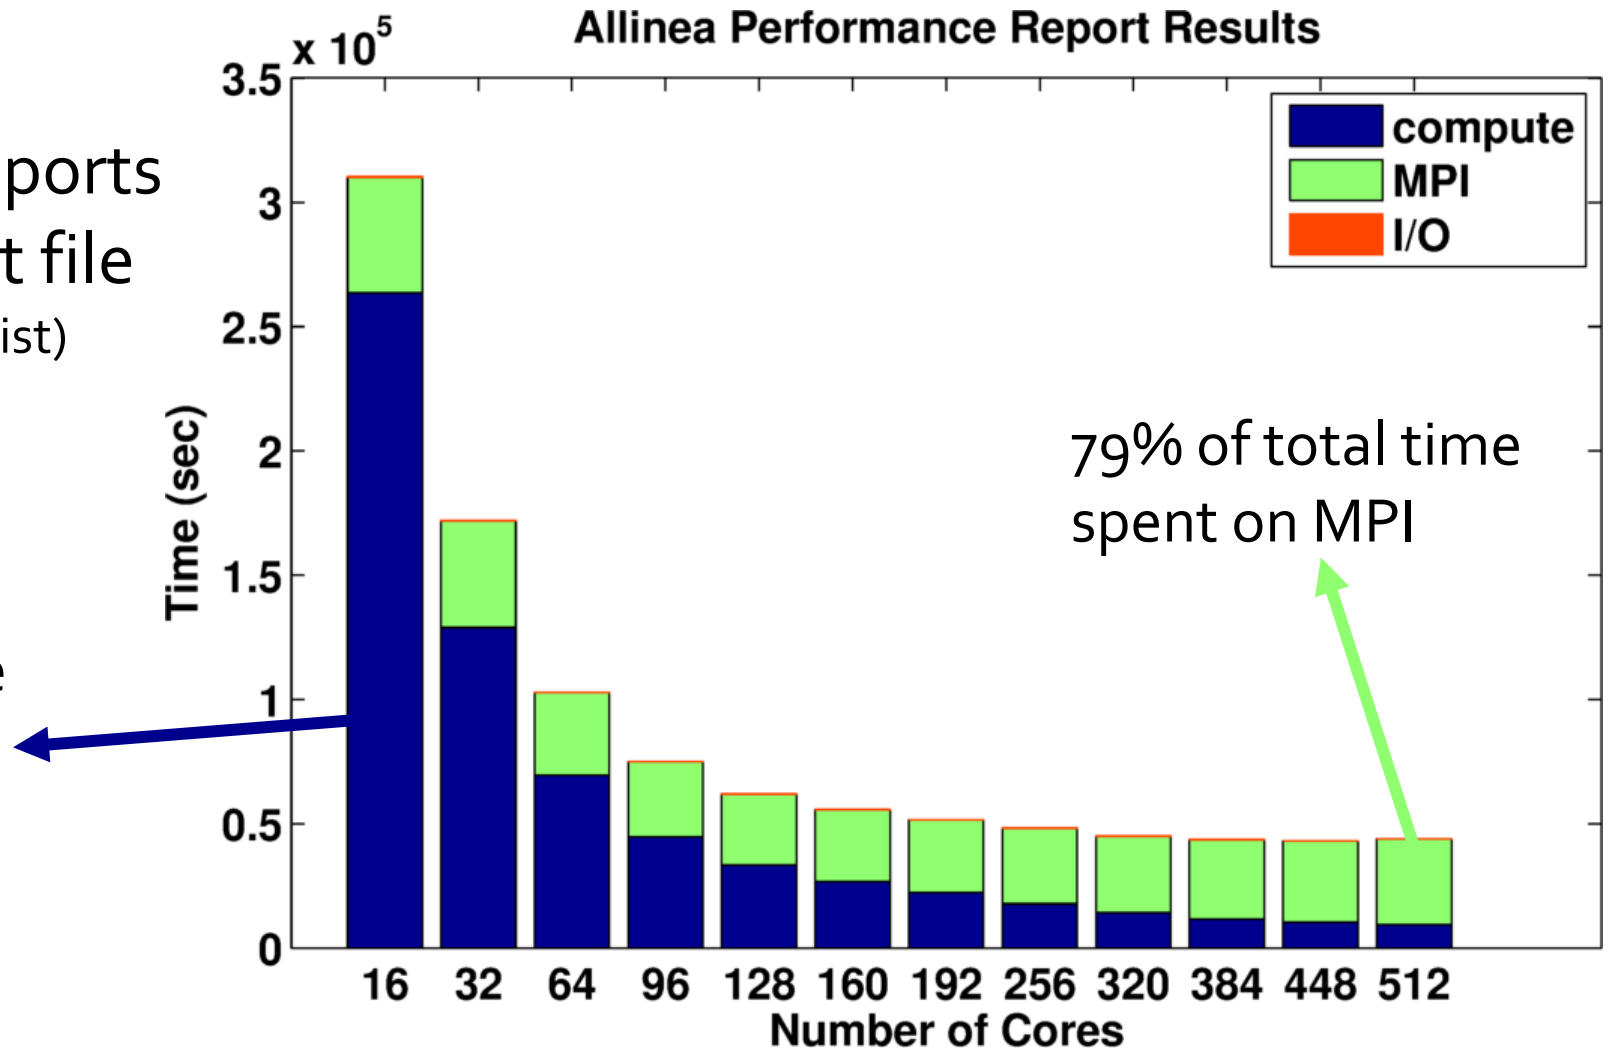


# Scalability Assessment (MPI Only)

4



- Allinea Performance Reports
- Separate netcdf output file  
(io\_form\_history=102 in WRF namelist)

87% of total time  
spent on  
computation



# Domain Decomposition (MPI only)

5




**Per Grid :  $1/4$  Computation and  $1/2$  MPI**

# Scalability Assessment (MPI Only)

6

- Allinea Performance Reports
- Separate netcdf output file  
(io\_form\_history=102 in WRF namelist)


87% of total time  
spent on  
computation



# Goal2. What does make WRF expensive?

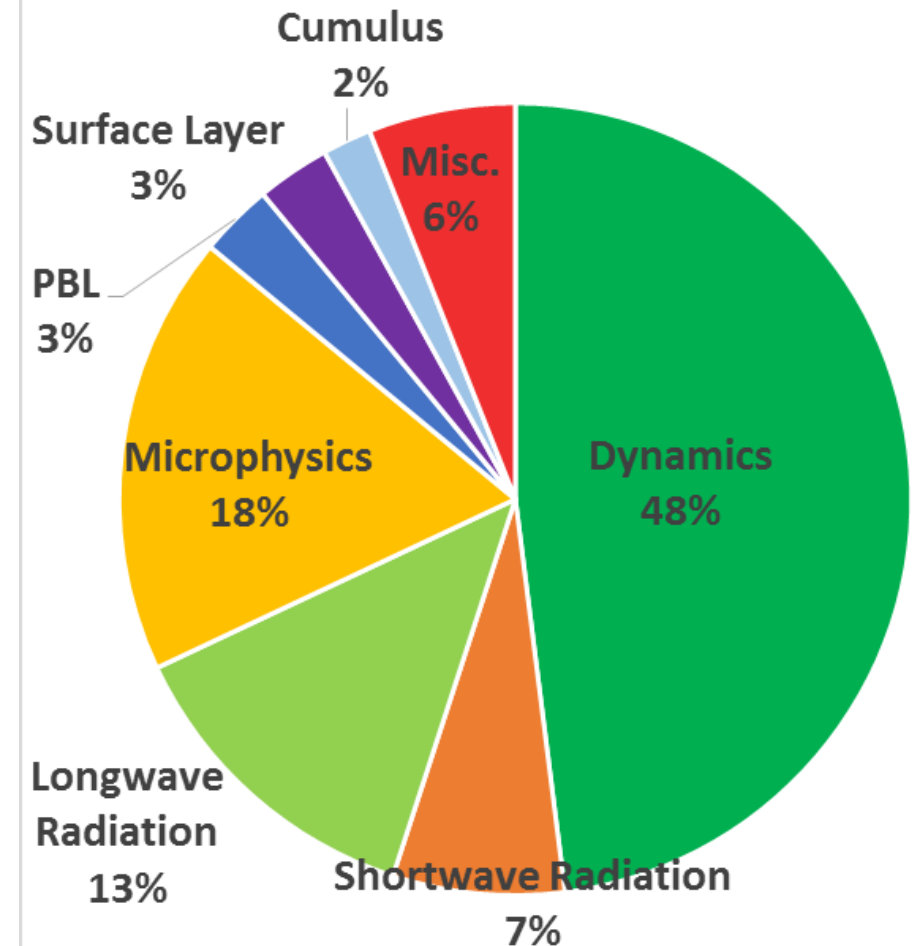
## Hurricane Sandy 4x4km case(500x500)

Domain Configuration



## Intel Vtune Amplifier XE Results

Time(%)




# Identifying hotspots and potential areas for improvement in WRF

## TAU/PAPI Tools Results

1- Positive Definite Delimiter (32 lines)

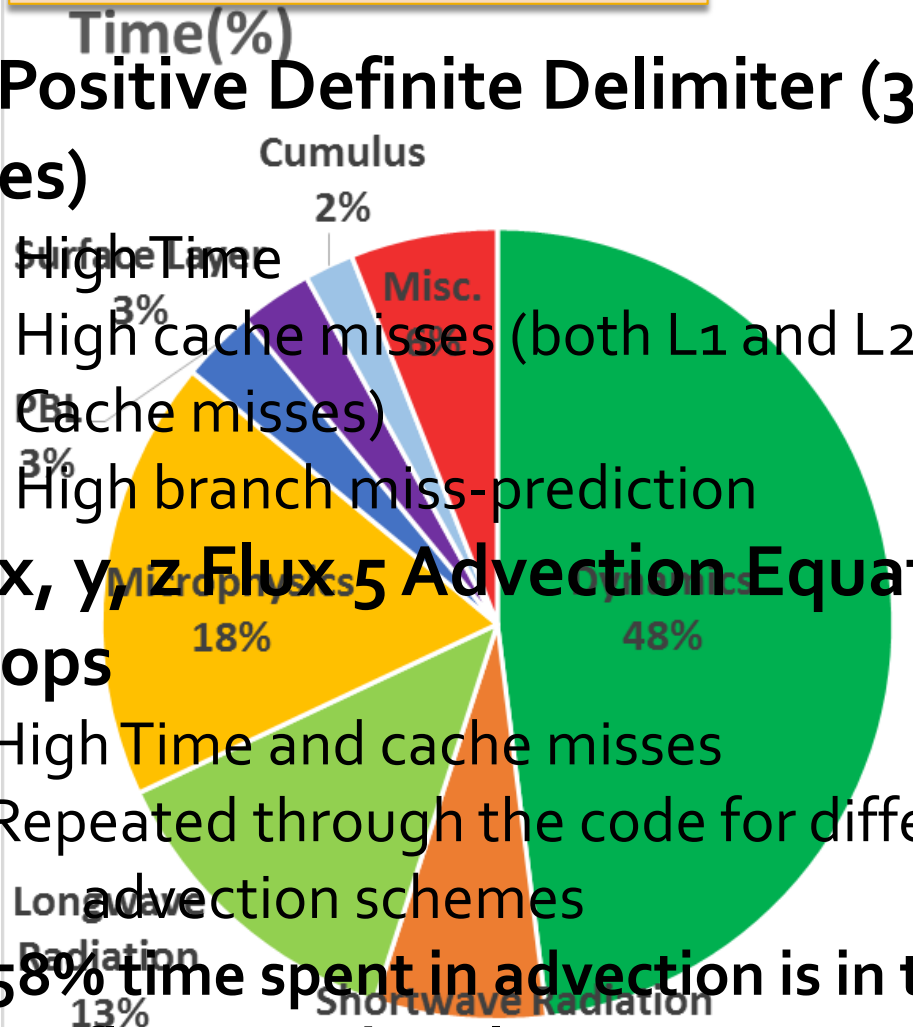
2- x, y, z Flux Advection Equation Loops

58% time spent in advection is in these flux equations loops



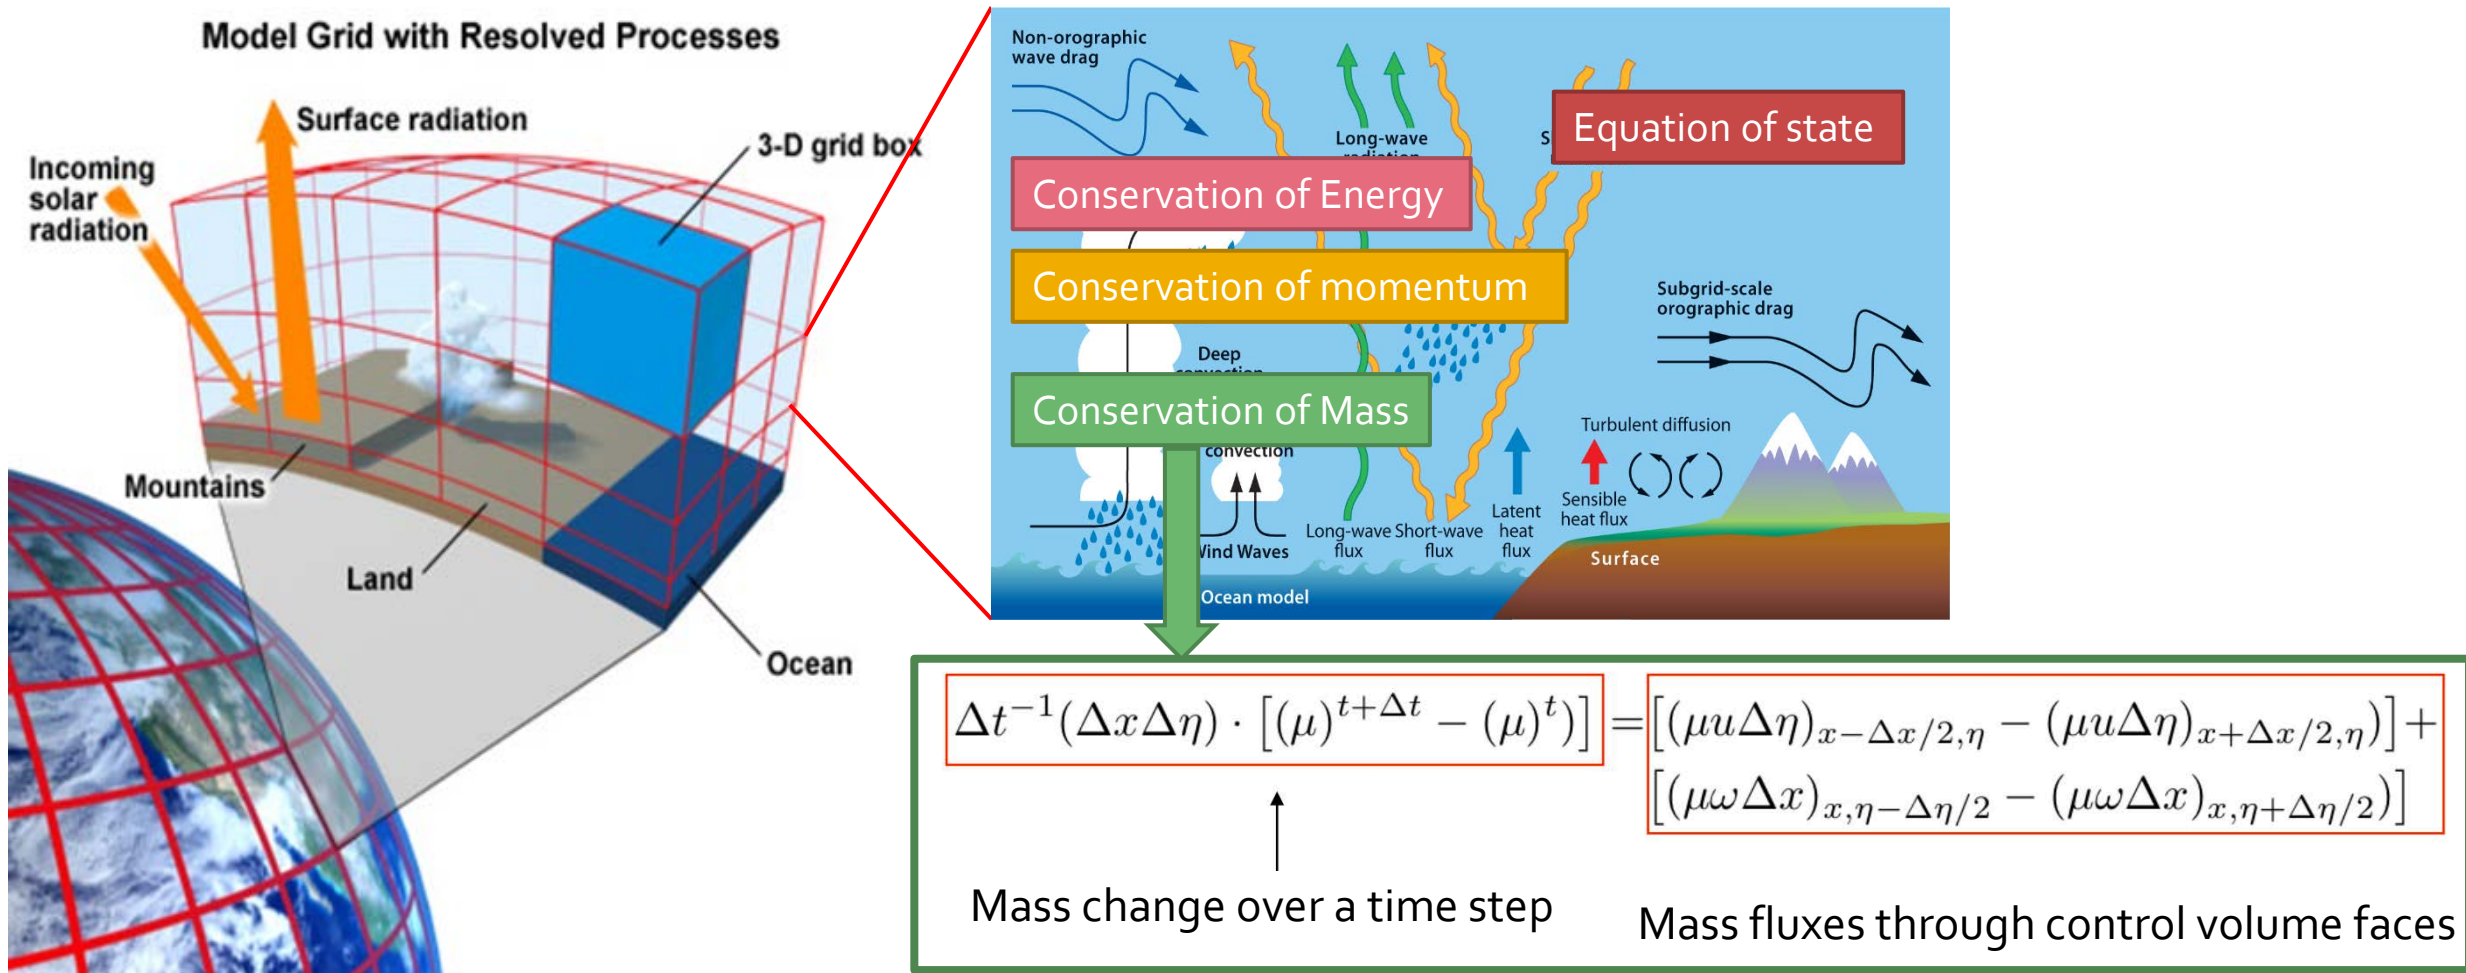
Detailed loop level high granularity analysis:

- 1- TAU tools
- 2- Allinea MAP





# Mass Conservation in the WRF Model



# Moisture transport in ARW

10





Figure from Skamarock and Dudhia 2012

- Until recently, many weather models did not conserve moisture because of the numerical challenges in advection schemes.  $\rightarrow$  high bias in precipitation
- WRF-ARW is conservative but not all of the advection schemes are.
- This introduces new masses to the system.

**Advection schemes can introduce both positive and negative errors particularly at sharp gradients.**

# Advection options in WRF

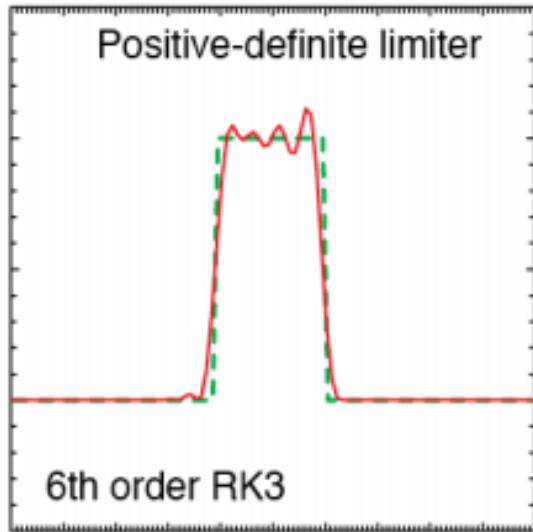
11



- High number of explicit IFs are causing high branch mispredictions

# Positive Define Advection Scheme

moist\_adv\_opt =1



## Hotspot

### Positive Definite Delimiter (32 lines)

High Time

High cache misses (both L1 and L2 Cache misses)

High branch miss-prediction

### Optimization Solution

#### Restructure and split the PD delimiter loop

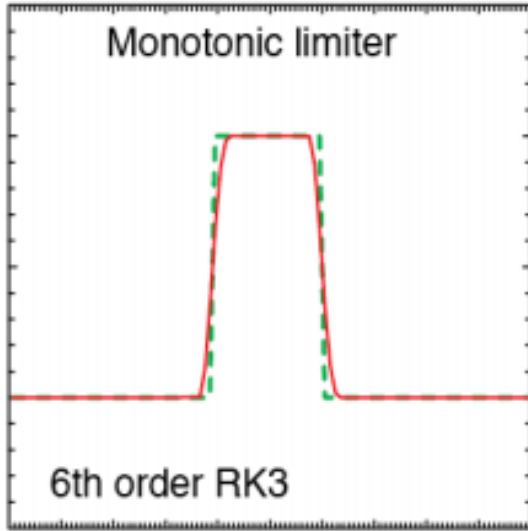
Increase vectorization

Reduce cache misses

Compiler	Optimization Flag	Loop Speed-up	Kernel Speed-up
Intel(v16.0.2)	-O3	100%	~17%
GNU (v6.1.0)	-Ofast	105%	~11%
PGI (v16.5)	-O3	35%	~4%

# Monotonic Advection Scheme

moist\_adv\_opt = 2



## Hotspot Monotonic Delimiter

High Time

High cache misses (both L1 and L2 Cache misses)

High branch miss-prediction

## Optimization Solution

Increase vectorization

Reduce cache misses

## Validation Tests:


- The results were bit-for-bit identical
- Regression tests with WTF (WRF Testing Framework).

Compiler	Optimization Flag	Whole Kernel Speed-up
Intel(v16.0.2)	-O3	~16%
GNU (v6.1.0)	-Ofast	~21%
PGI (v16.5)	-O3	~9%

# Optimization strategy

- **Expose more parallelism**
  - Push columns inside
  - Facilitates better vectorization
  - Facilitates cache blocking
- **Force vectorization** using SIMD & vector align directives
- **Replace:**
  - divisions with reciprocal multiplications
  - repeated indexed access with copies having coalesced access
  - assumed-shaped arrays (e.g. (:,:)) with proper declarations)
- **Eliminate:**
  - unnecessary code, variables, initializations or copies

# Positive Define Advection Scheme



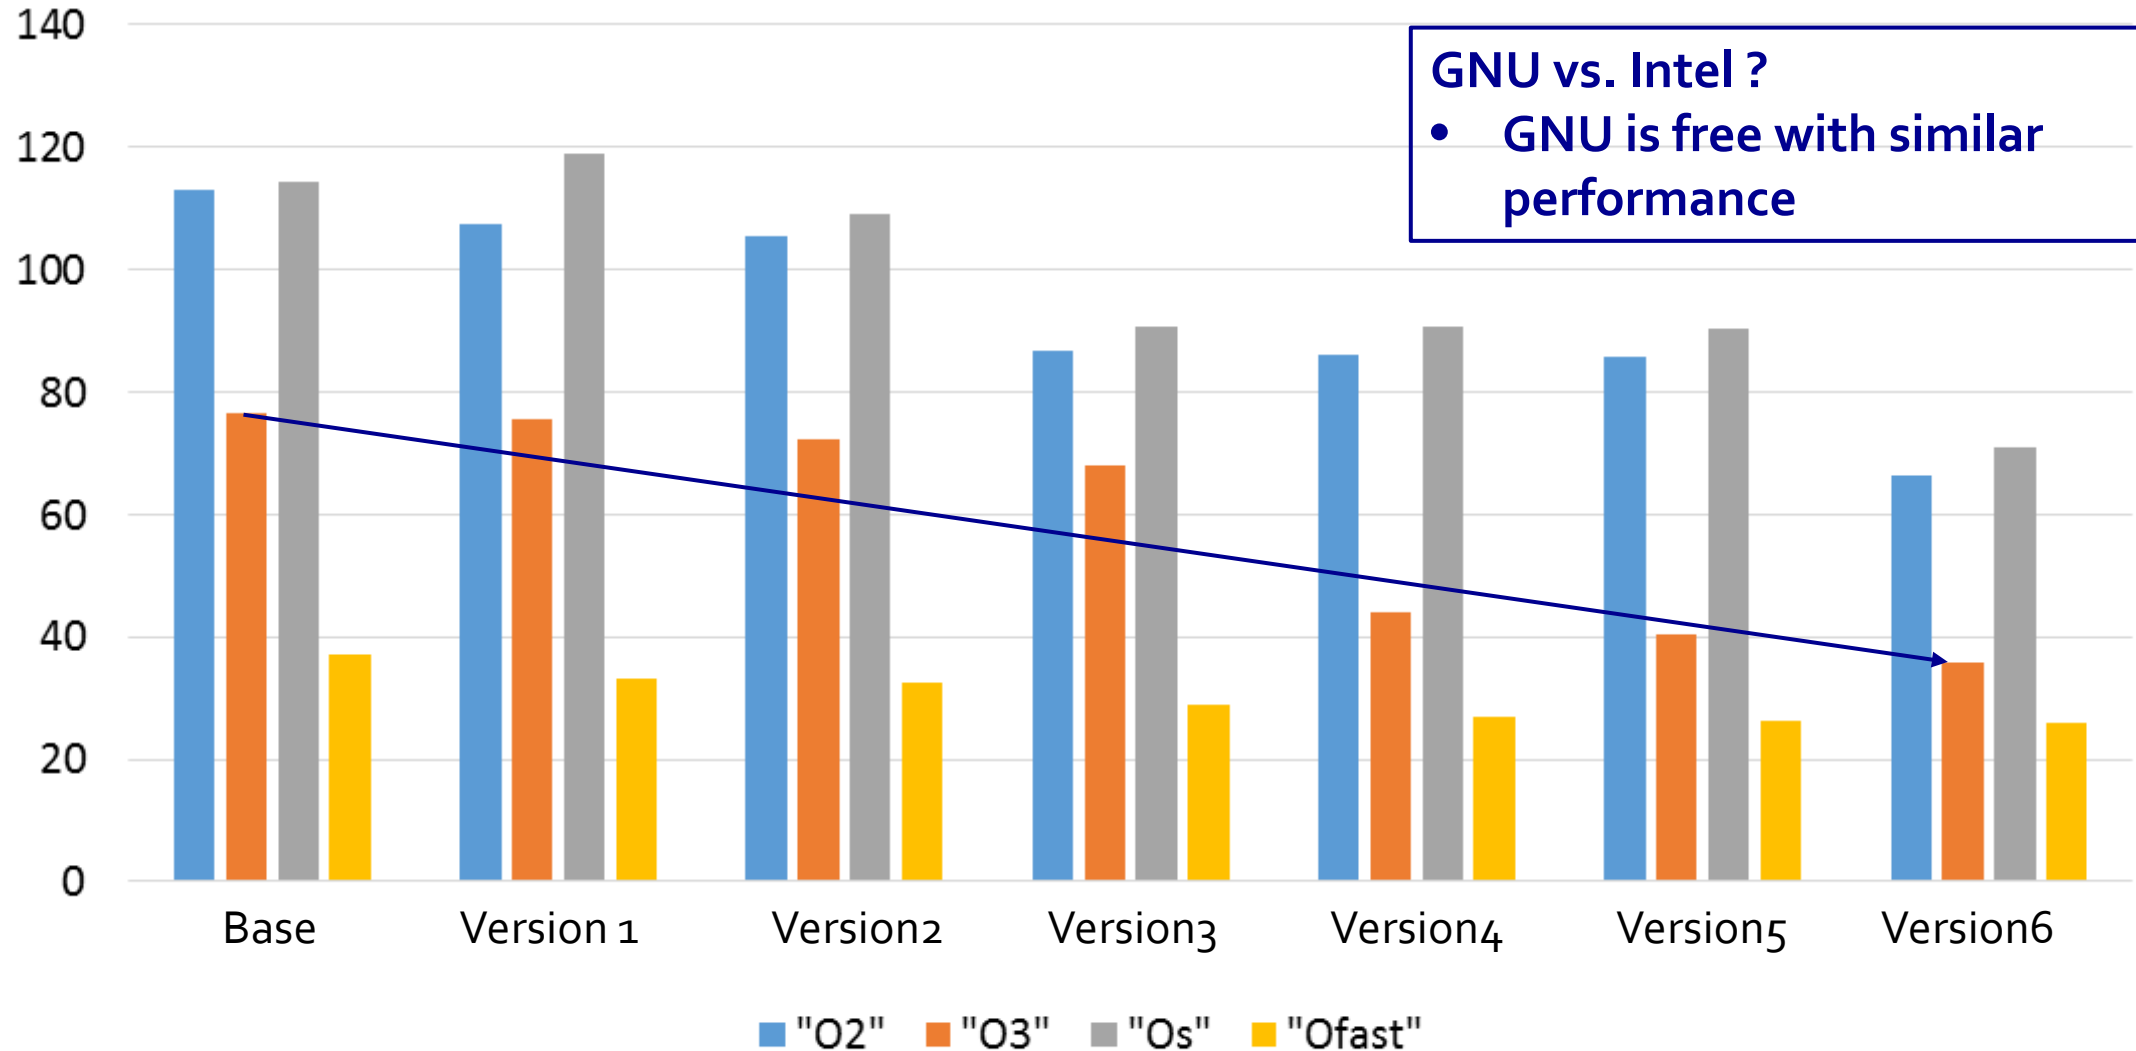
# Positive Define Advection Scheme

GNU-6.1.0

>1.8x speed-up with GNU compilers


GNU vs. Intel ?

- GNU is free with similar performance





# Positive Define Advection Scheme



# Conclusion

18

- **Profiling:** WRF with Intel Vtune XE, TAU tools, and Allinea MAP for identifying the hotspots of WRF
  - ▣ Dynamics is identified as the most expensive part of ARW
- **Optimizing:** the identified hotspots of different advection schemes for Intel, GNU , and PGI compilers
  - ▣ Significant speed-up of the advection schemes
- **Validating:** the results of the advection kernel
- **Integration:** The changes to WRF advection schemes are approved by the WRF committee and **integrated** in the main WRF repository.

# Ongoing and Future Work

19

## **Performance Improvement of advection modules**

- Reducing memory footprint by decreasing the number of temporary variables
- Increasing the advection module vectorization
- Analysis of hardware counters to fix branch mispredictions and cache misses

**Future pull requests will provide up to 2.5x speed-up for advection schemes.**

# Acknowledgements

20

- ❖ Davide Del Vento
- ❖ Rich Loft
- ❖ Srinath Vadlamani
- ❖ Dave Gill
- ❖ Dave Hart
- ❖ Siddhartha Ghosh
- ❖ Greg Carmichael

**XSEDE**

Extreme Science and Engineering  
Discovery Environment

**TACC**

TEXAS ADVANCED  
COMPUTING CENTER



**CISL** Computational & Information Systems Laboratory