KIAPS
KOREA INSTITUTE OF
ATMOSPHERIC PREDICTION SYSTEMS

Improving the Performance of the Global Atmospheric Model with OpenMP

Kwangjae Sung, Tae-Jin Oh and Junghan Kim

Oct 4, 2016 / OpenMPCon 2016

www.kiaps.org

Outline

- KIAPS & KIAPS Integrated Model (KIM) Introduction
- KIM Performance with OpenMP

KIAPS Introduction

- Brief Introduction
  - KIAPS - **K**orea **I**nstitute of **A**tmospheric **P**rediction System
  - 9 Year project, approx. $100M funded by South Korean Government
- Vision & Goals
  - To develop the next generation global NWP system optimized to the topographic & meteorological features of Korea
  - To reduce the economic loss caused by natural disasters and enhance productivity of industrial sector
  - To build science & technology capacity that stimulates NWP research

KIAPS Roadmap

step1

Foundation and basic research

- Establish the foundation of KIAPS research and development environment
- Efforts on laying out future model development – model dynamics, physics and data assimilation

2011-2013

step2

Release beta version

- Complete developing major model components based on KIAPS own research
- Release the KIAPS beta version model
- conduct semi realtime experiments with KIAPS beta version and evaluate it skills to the KMA operational model

2014-2016

step3

Finalize the operational system

- Finalize the KIAPS operational version 1.0 and run on semi-operational basis
- Stabilize the model system by further diagnostics and verifications
- Release the KIAPS system to international users

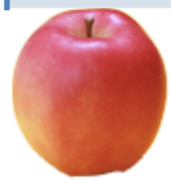

2017-2019

Supercomputing Trends

- Roughly x1000 computation power increase in 10 years
- Clock speed → Core count (heat problem)
- Scalability will become one of the key issues in NWP modeling

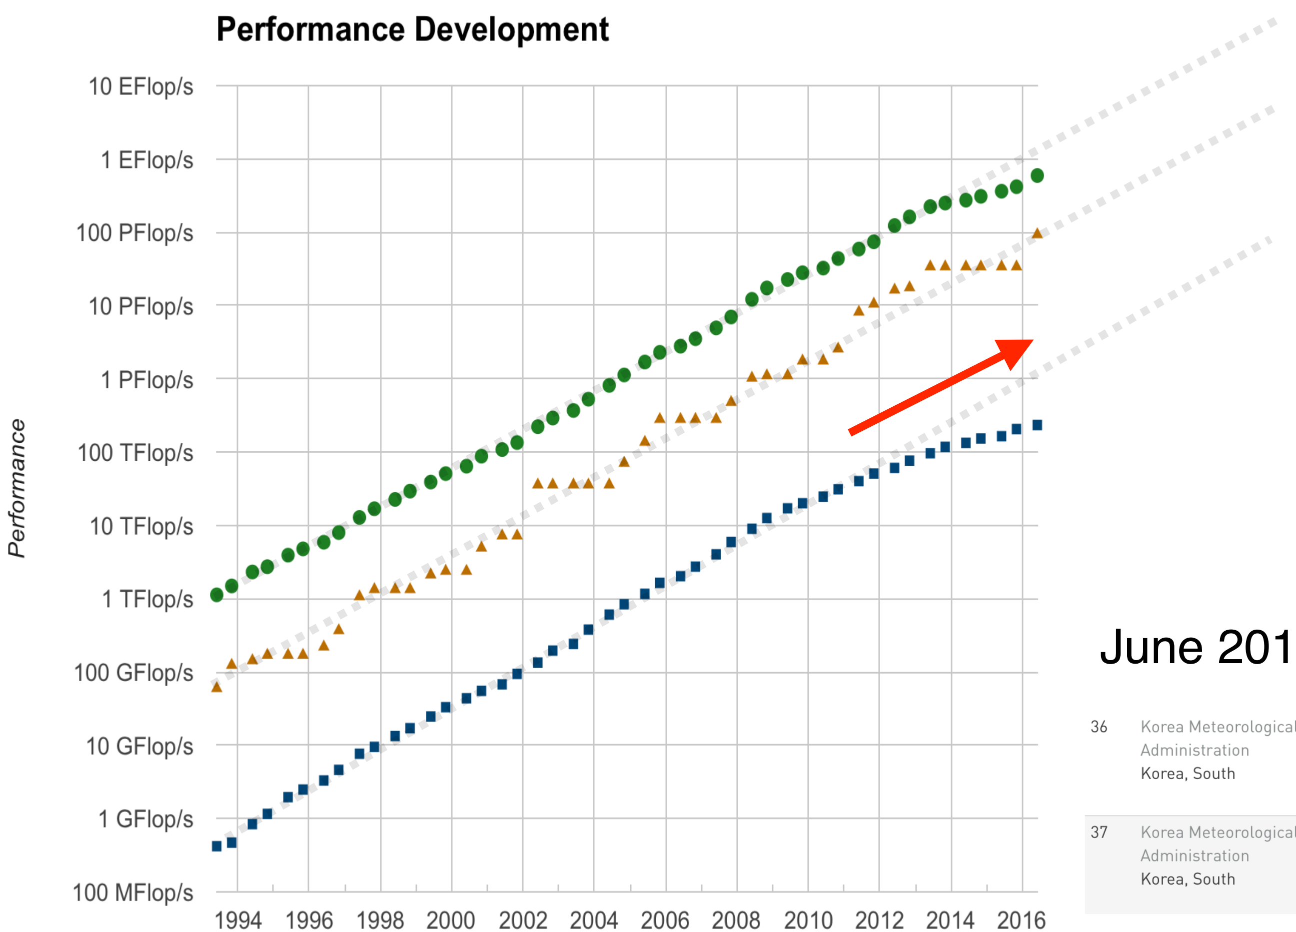

June 2016, http://top500.org

36	Korea Meteorological Administration Korea, South	**Miri** - Cray XC40, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect Cray Inc.	69,600	2,395.7	2,895.4
37	Korea Meteorological Administration Korea, South	**Nuri** - Cray XC40, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect Cray Inc.	69,600	2,395.7	2,895.4

UM Scalability

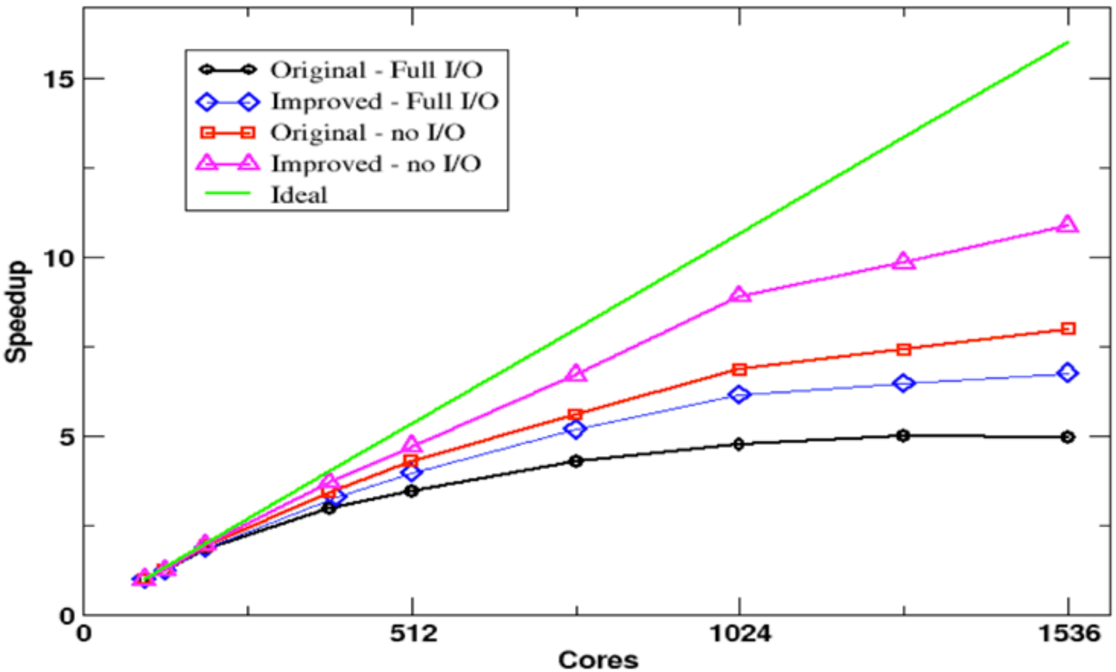

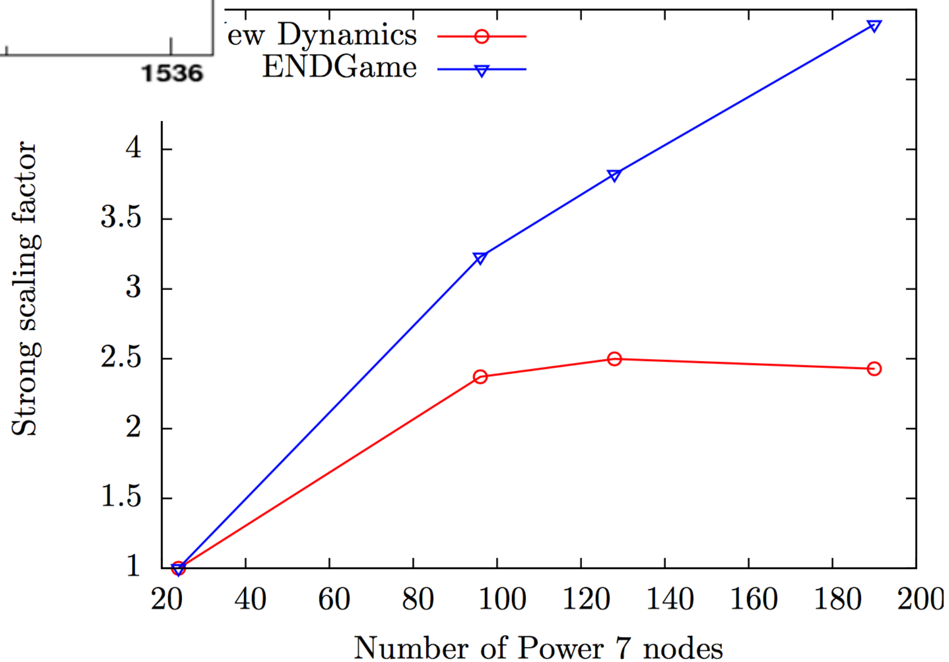

Strong Scaling plot for ENDGame and New Dynamics forecast at N768 resolution on the IBM Power 7. Baseline forecast on 24 nodes. 1 IBM node contains 32 processors.

Lat-Lon grid and Alternatives

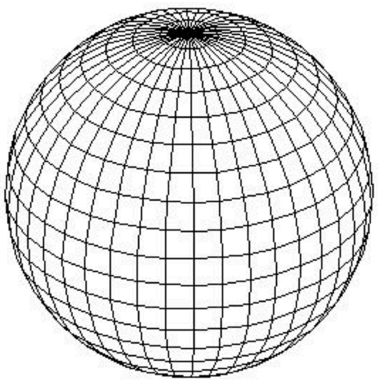

- Pros:
  - Orthogonal
  - Easy implementation of high order discretization
- Cons:
  - Anisotropic
  - Resolution variance
  - Pole Singularity
  - Poor scalability

- Pros:
  - Isotropic
  - Highly uniform resolution
- Cons:
  - Non orthogonal
  - Difficult to implement high order discretization
  - Difficult local refinement
  - 12 special cells (pentagon)

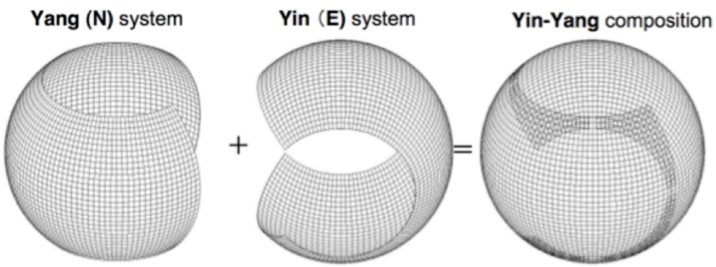

Provided by Dr. Kageyama, ESC, who is the developer of the Yin-Yang grid

- Pros: Isotropic, uniform resolution, orthogonal, local refinement is possible
- Cons: Overlaps need special treatment, local refinements at boundaries would need special treatment

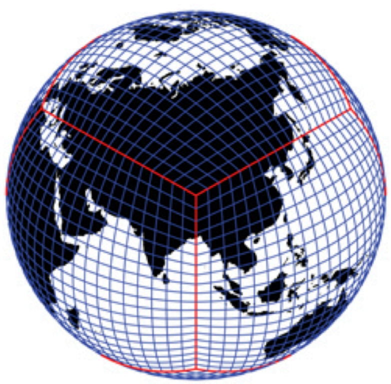

- Pros:
  - Isotropic
  - Relatively uniform resolution
  - Local refinement possible
- Cons
  - 8 special corner point
  - Plane boundaries need special treatment

Next Generation NWP

- Global, Fully Compressible, Nonhydrostatic Model
  - **High-resolution** Cloud Resolving Earth System Model
  - Systematic Global + Regional high resolution simulation
  - Model that is capable of simulating across scales (Temporal & Spatial)
- **Scalable** on Systems of CPU O(1E5) and beyond
  - No pole singularity on the grid structure
  - Local numerical procedures (minimum communication)
- Mass conservation
- Shape-preservation, positive-definite, monotonicity, non-oscillatory for required scalars
- High order accuracy
- Computationally **Efficient** (i.e., that satisfies operational cutoff time)
- Adaptive Mesh Refinement (AMR) capability (optional)
- **Readiness for unknown target architecture (CPU, GPU, Xeon Phi, ?)**

2012 DyCore Model Inter-comparison Project

	Origin	Eqns.	Grid Structure	Horizontal	Vertical
CAM-FV	NCAR	H	lat-lon	C D	Lin-Rood
CAM-SE	NCAR	H	Cubed S	SE	sigma-p L
Dynamico	France	H	Icos(hex)	C	sigma-p L
ENDGame	UKMO	N	lat-lon	C	sigma-z CP
FIM	NOAA	H	Icos (hex)	A	theta
FV3	GFDL	H	Cubed S	C	Lin-Rood
GEM	Env Ca	H	lat-lon	C	sigma-p CP
GEM-YY	Env Ca	H	Yin-Yang	C	sigma-p CP
IFS	ECMWF	H/N	Gaussian	SH	FEM
ICON-IAP	IAP	N	Icos (hex)	C	sigma-z L
ICON	MPI-DWD	H/N	Icos (tri)	C	H:s-p , N:s-z, L
MCORE	U Michigan	N	Cubed S	A	
MPAS	NCAR	H/N	SCVT	C	H:s-p , N:s-z, L
NICAM	Japan	N	Icos (hex)	A	
NIM	NOAA	N	Icos (hex)	A	sigma-z L
OLAM	U Miami	H	Icos (tri)	A	cut-cell z
PUMA	U Hamburg	H	Gaussian	SH	sigma-p
UZIM	CSU	N	Icos (hex)	Z	Arakawa Konor

KIAPS DyCore Options

- Option 1
  - Icosahedral / SCVT
  - Finite Volume
  - NCAR MPAS, NOAA FIM/NIM, DWD/MPI-M ICON, CCSR/JAMSTEC NICAM
- **Option 2**
  - Cubed sphere
  - Spectral Element (SE, CG for Continuous Galerkin) / Discontinuous Galerkin
  - NCAR CAM-SE, NPS NUMA
- Option 3
  - Yin-Yang grid
  - Finite Volume / Semi-Lagrangian
  - CMC-GEM

Spectral Element Model Scalability

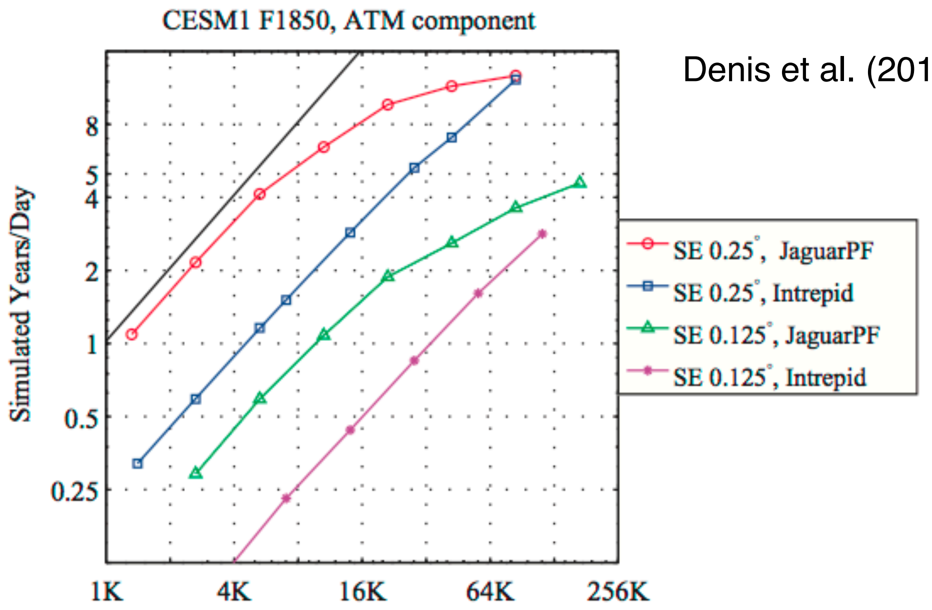

Denis et al. (2012) CAM-SE, CESM-atm component with full physics

Kelly and Giraldo (2012) NUMA

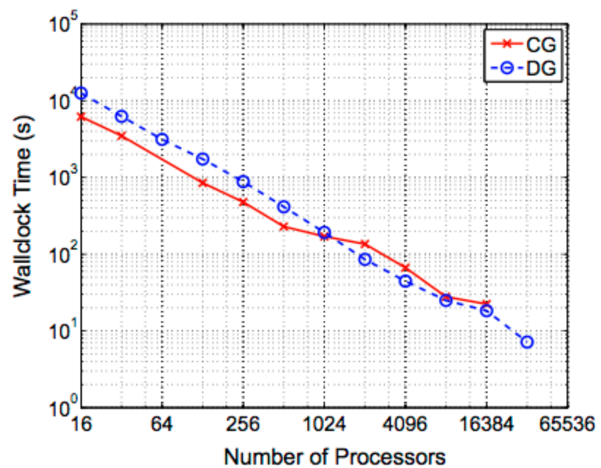

(a) wallclock time

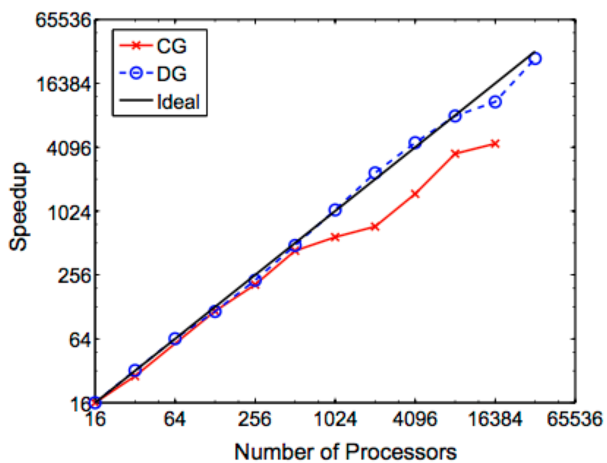

(b) speedup

Spectral Element Communication Footprint

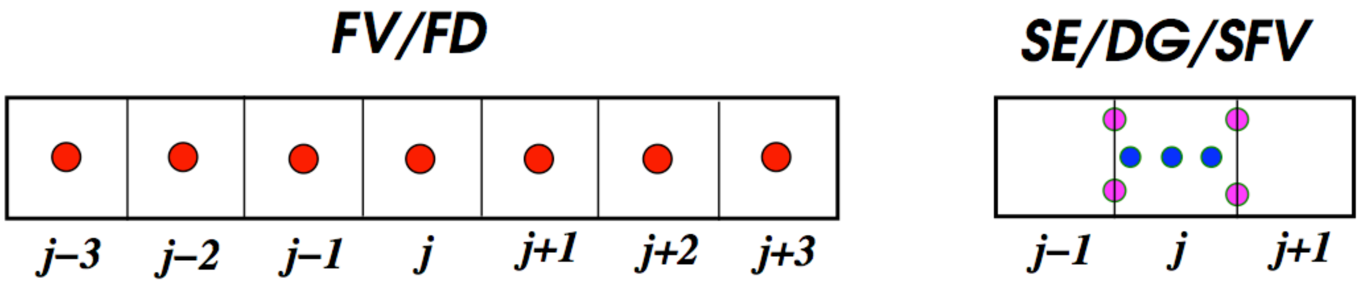

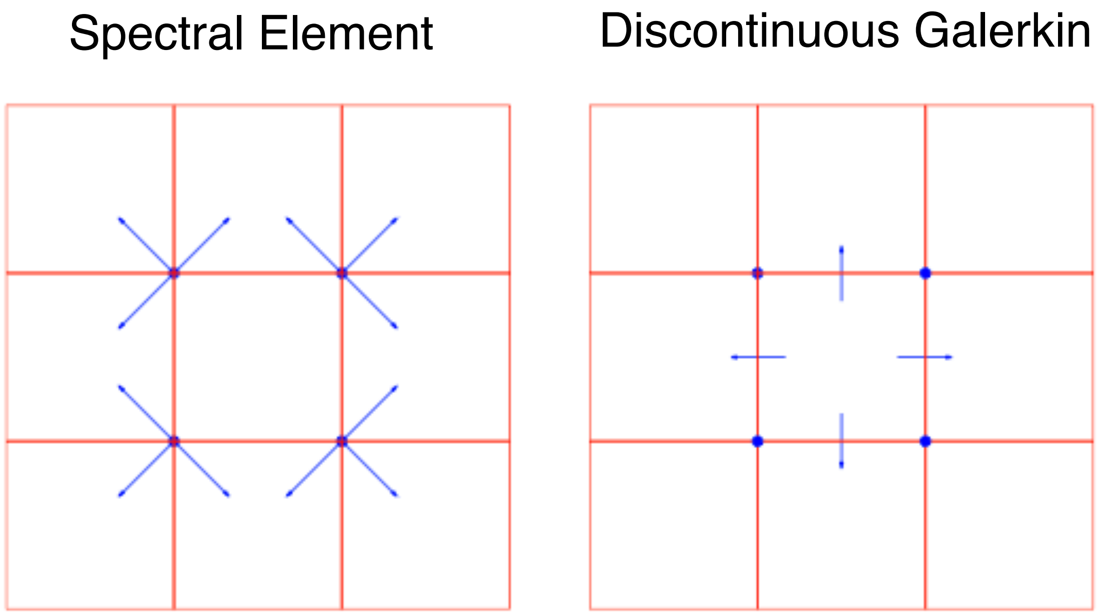

KIAPS Integrated Model (KIM)

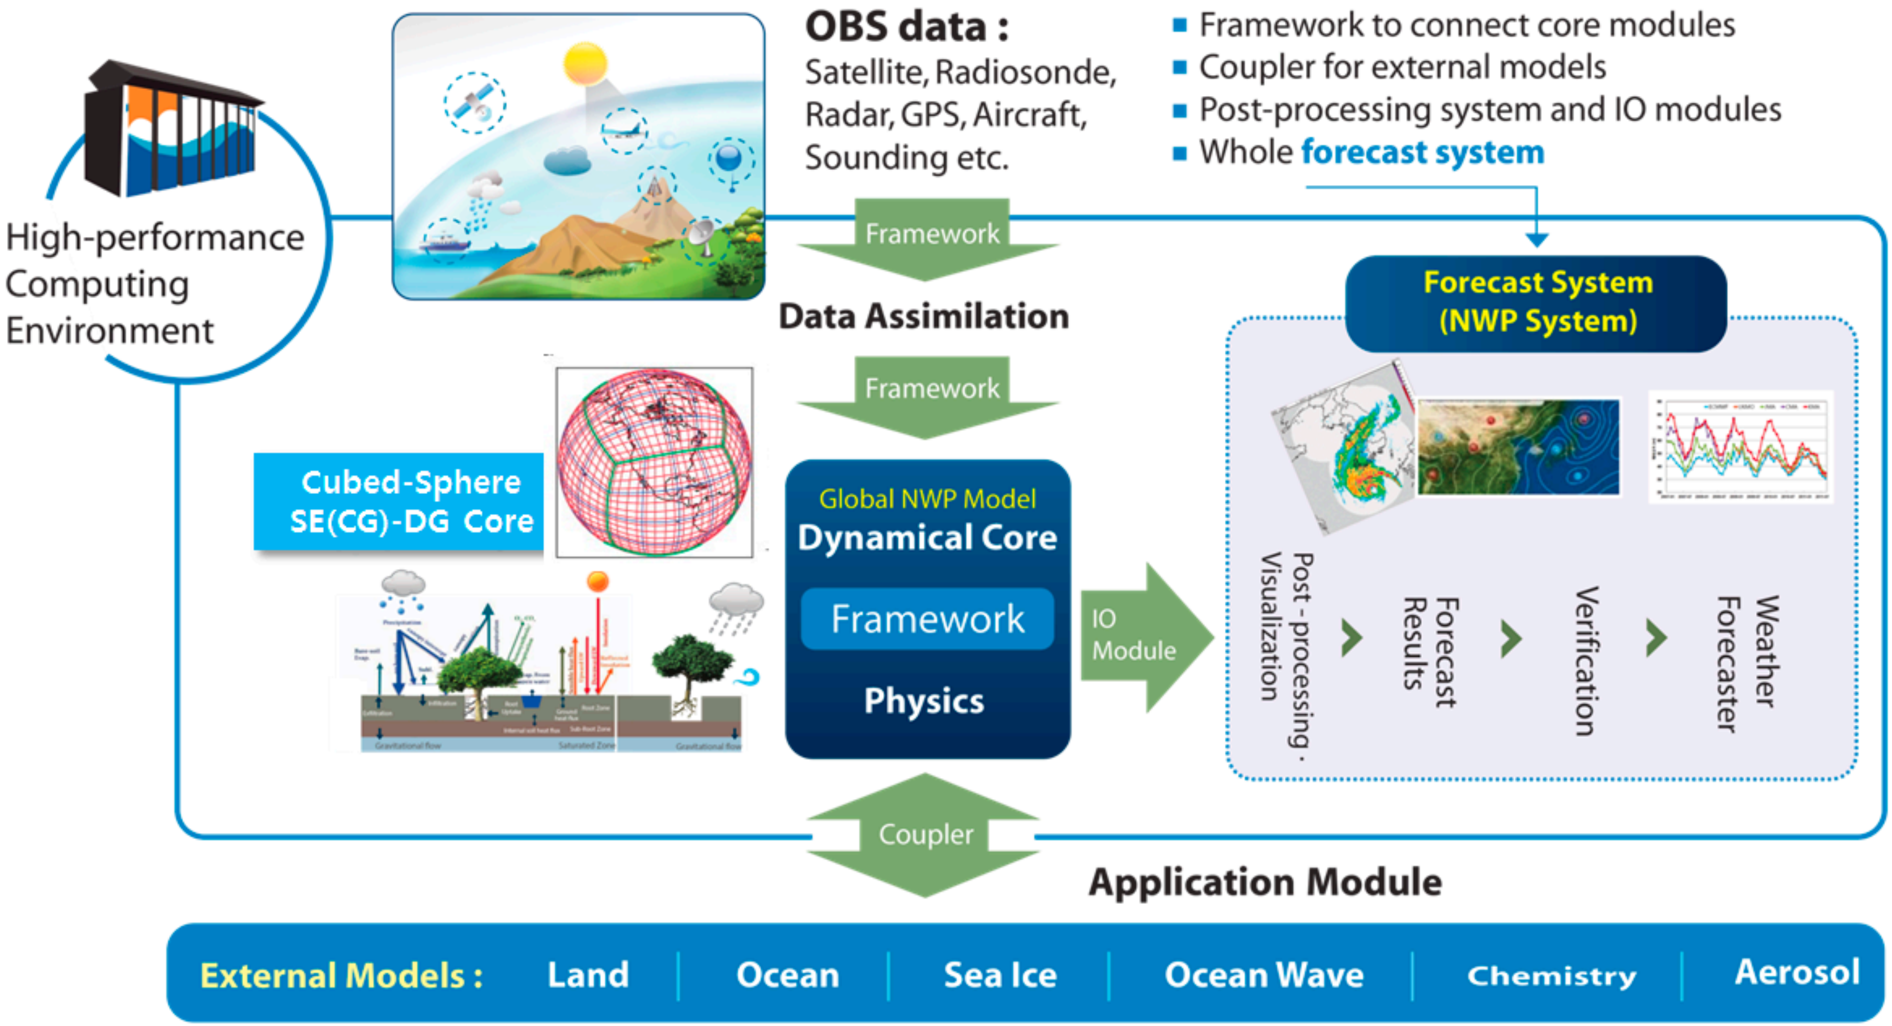

KIM Development History

		KIM1.0	KIM2.0	KIM2.1
Dynamical core	Equation	Hydrostatic	Nonhydrostatic	—
	Horizontal discretization	Spectral element method	—	—
Physics	Shortwave rad.	RRTMG	—	—
	Longwave rad.	RRTMG	—	—
	Land surface	NOAH V2.5	NOAH V3.0	+ 3-layer sea-ice model
	Ocean surface layer	N/A	Kim and Hong (2010)	—
	Vertical diffusion	YSU (+stable BL)	+ top-down mixing	—
	Gravity wave drag	O: McFarlane (1987) C: Warner and McIntyre (2001)	O: Alpert et al. (1996), Hong et al. (2009) C: Chun and Baik (1998), Jeon et al. (2010)	+ Rayleigh diffusion (O)
	Deep conv.	SAS (Han and Pan 2011)	+ scale aware (Lim et al. 2014) + updates in 2015	+ Minor bug fixed
	Shallow conv.	Han and Pan (2011)	Hong et al. (2013)	—
	Microphysics	WSM6 (Hong and Lim 2006)	WSM5 (Hong et al. 2004)	+ Minor bug fixed
	Cloudiness	Wilson and Gregory (2003)	Prognostic scheme (Park et al.)	+ revised CPS condensate

Mountain Wave

ENDGame

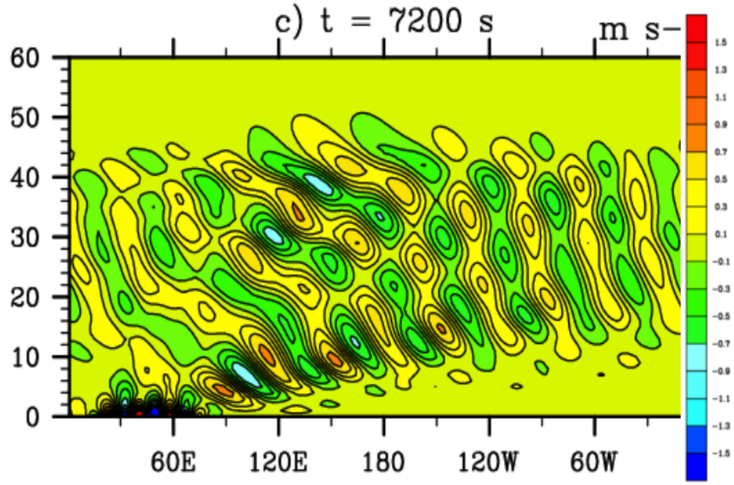

NIM

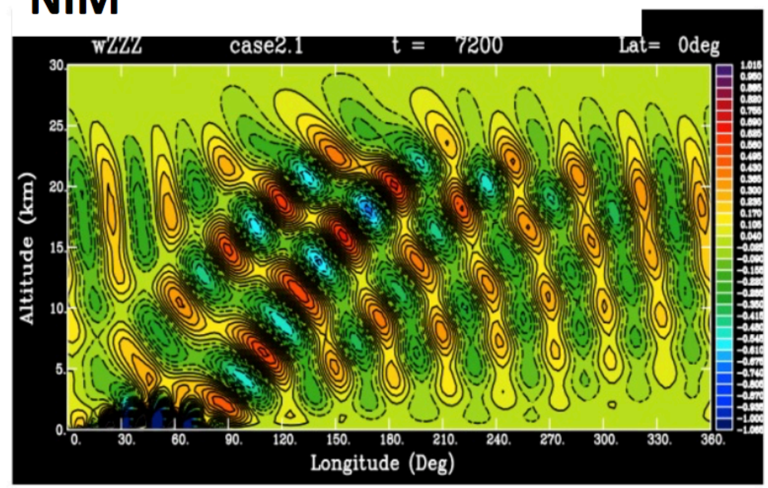

MPAS

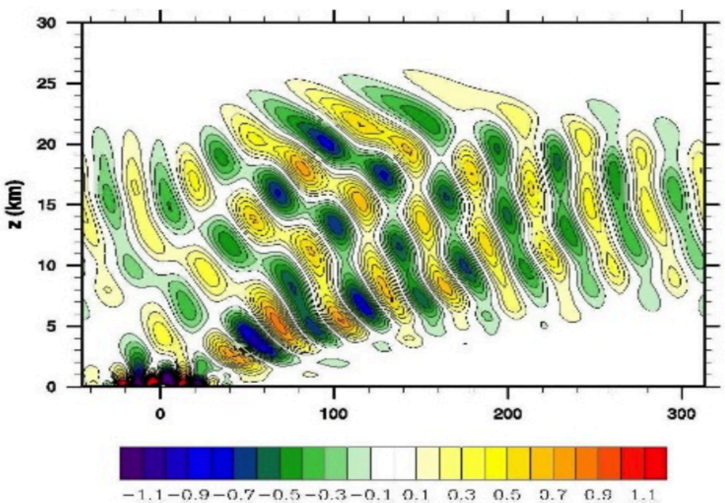

KIM

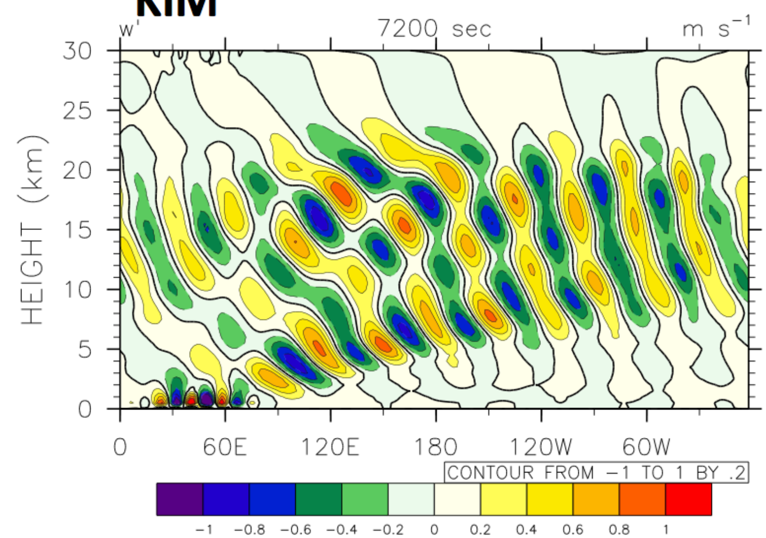

Typhoon Bolaven / Physics Verification

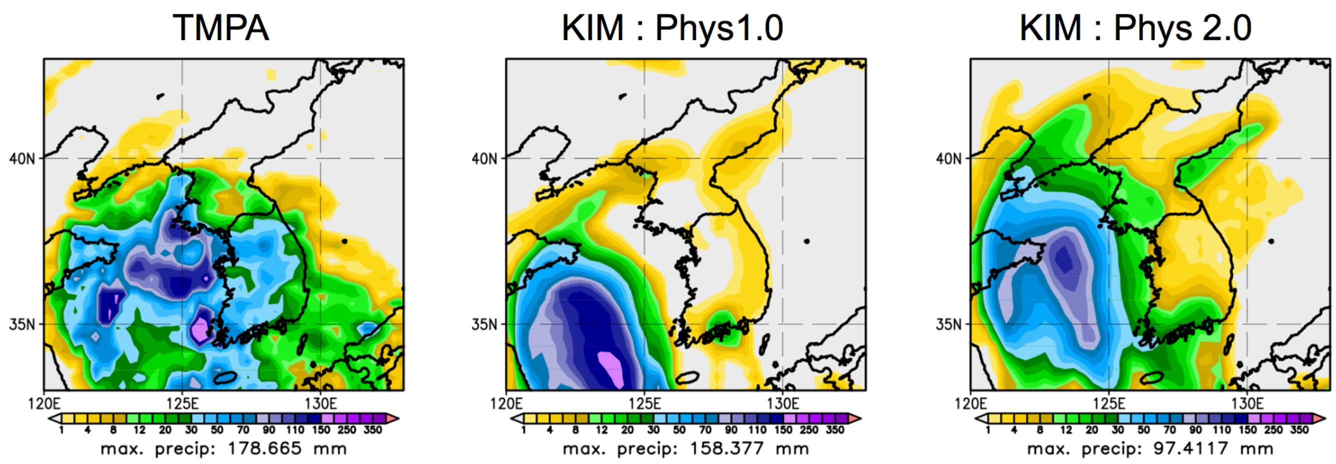

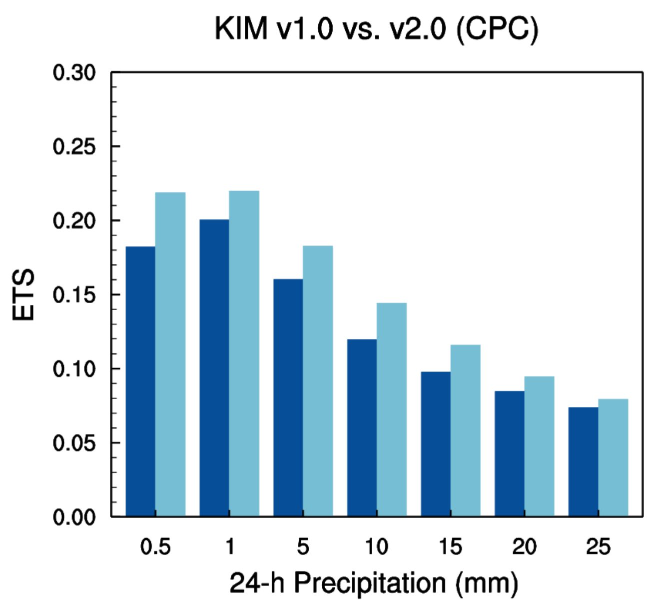

[against CPC observation]

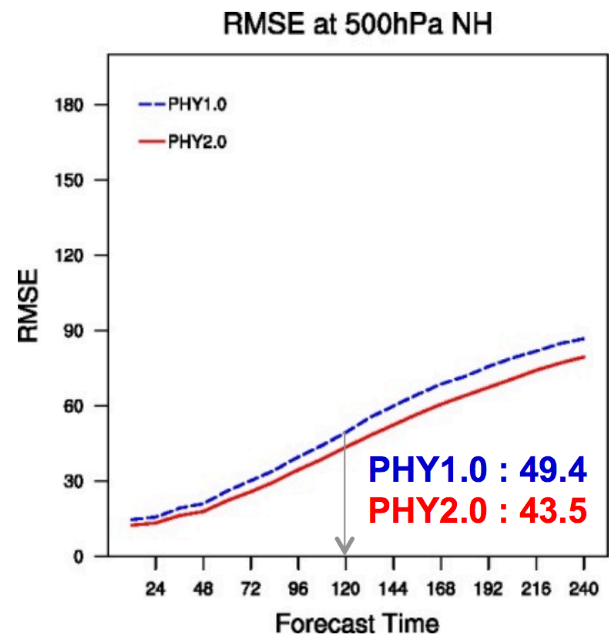

[against FNL analysis]

Semi-Real Time Operation Status

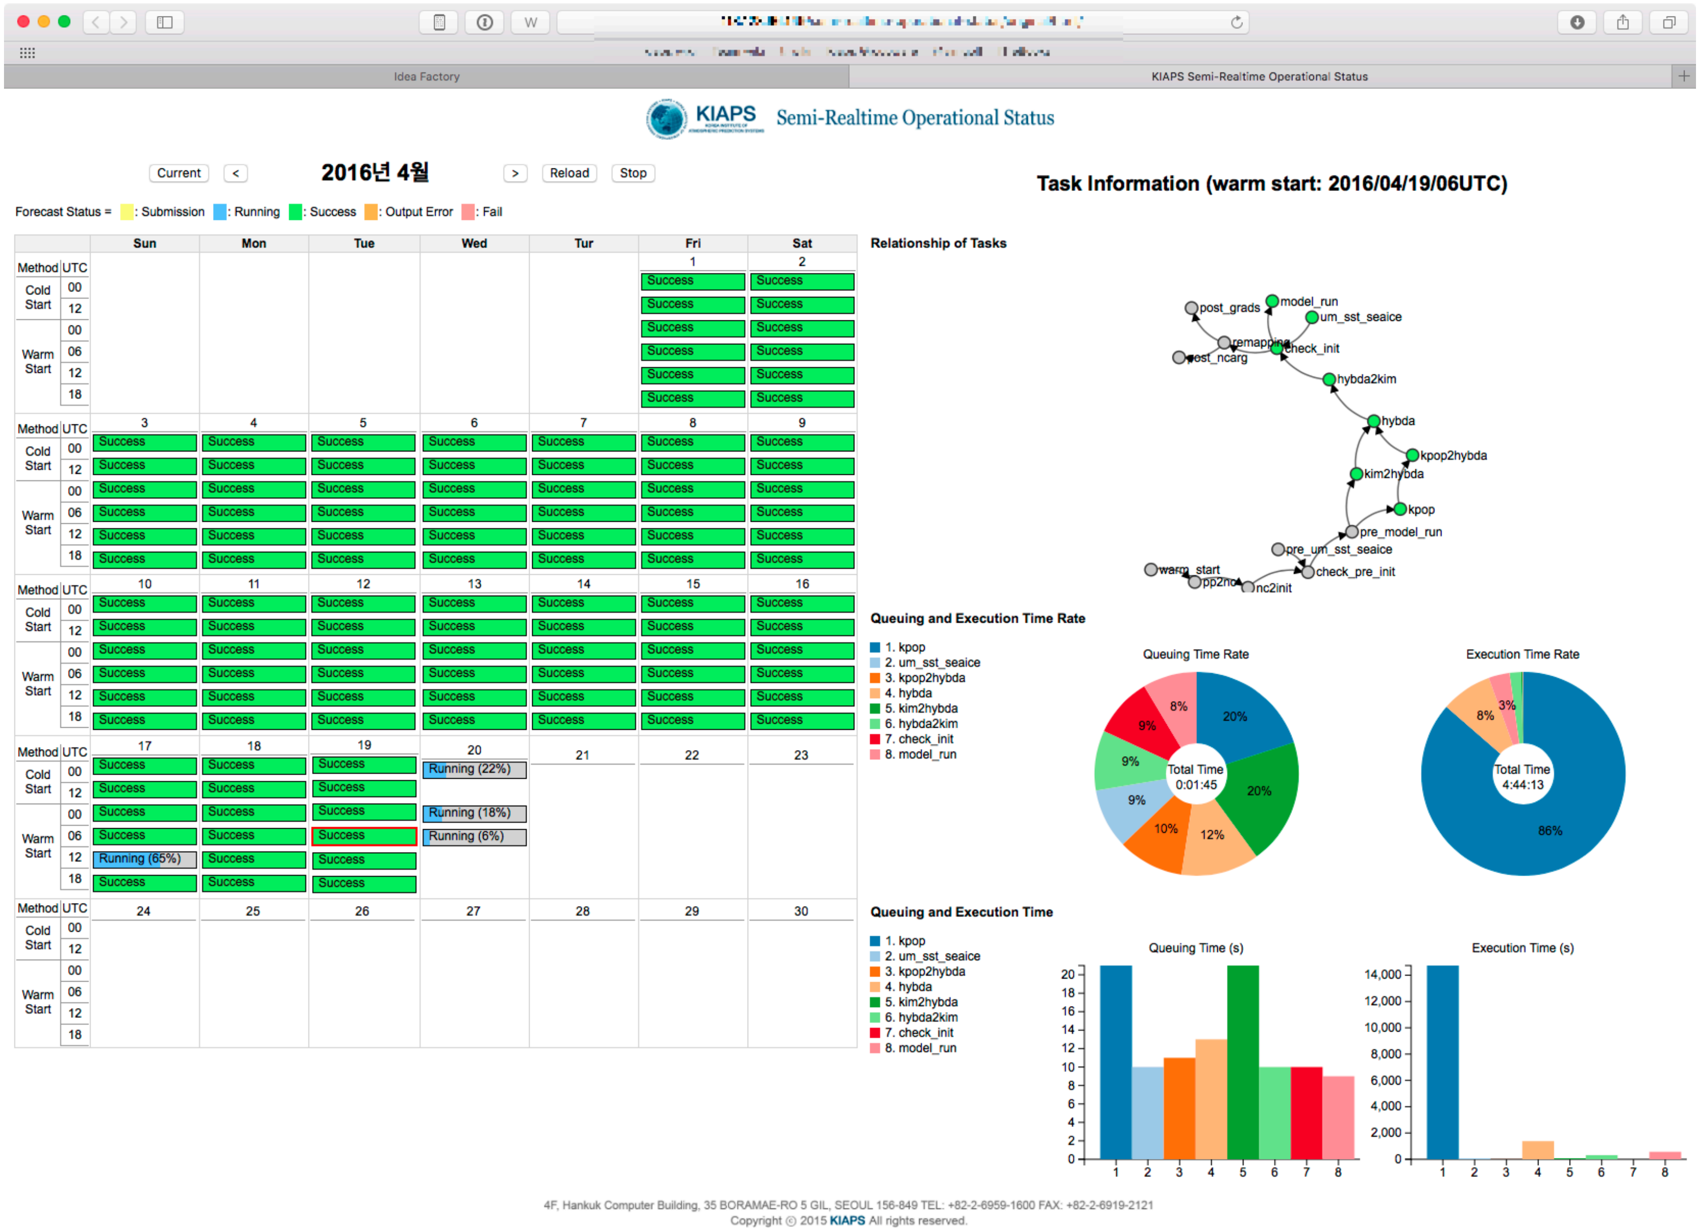

Megi Track Prediction Composite

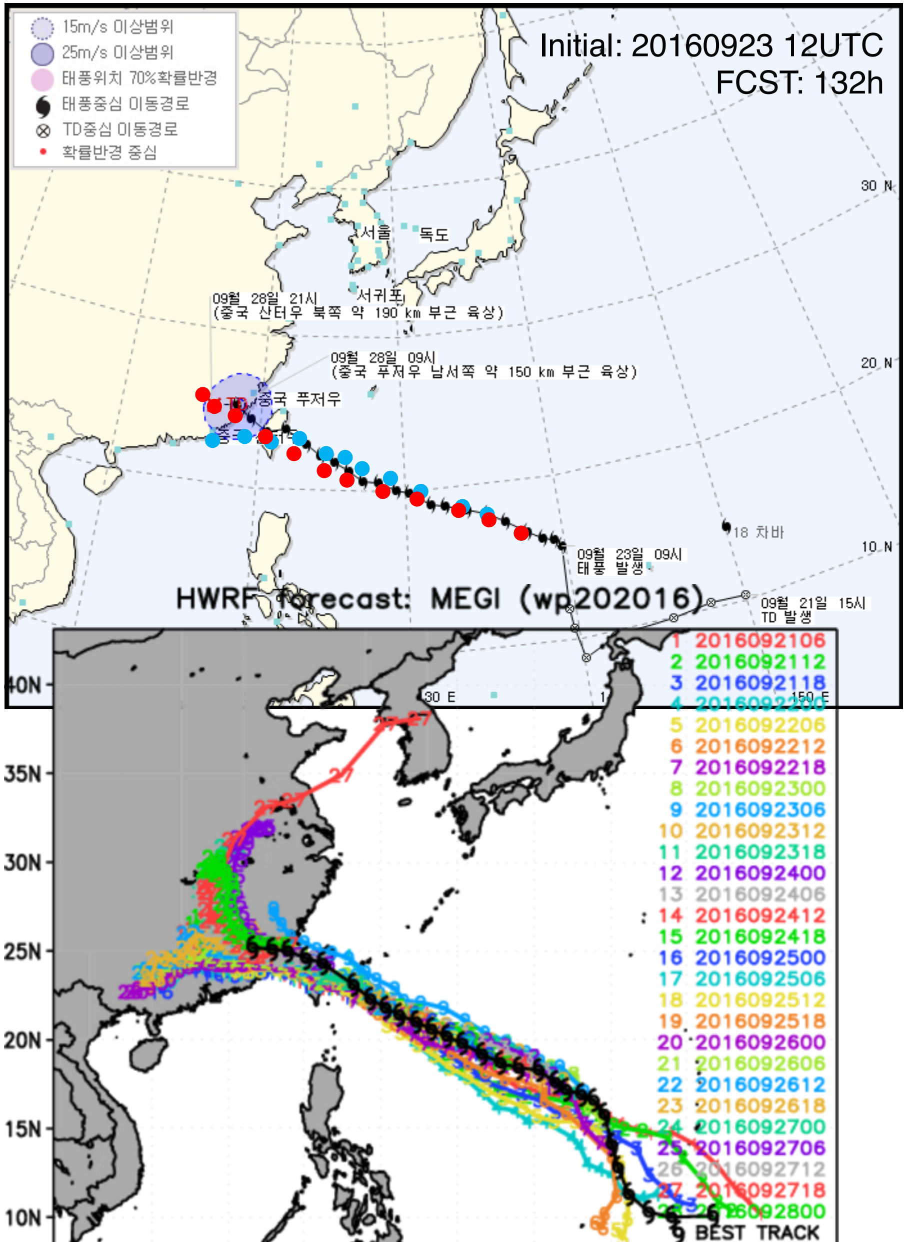

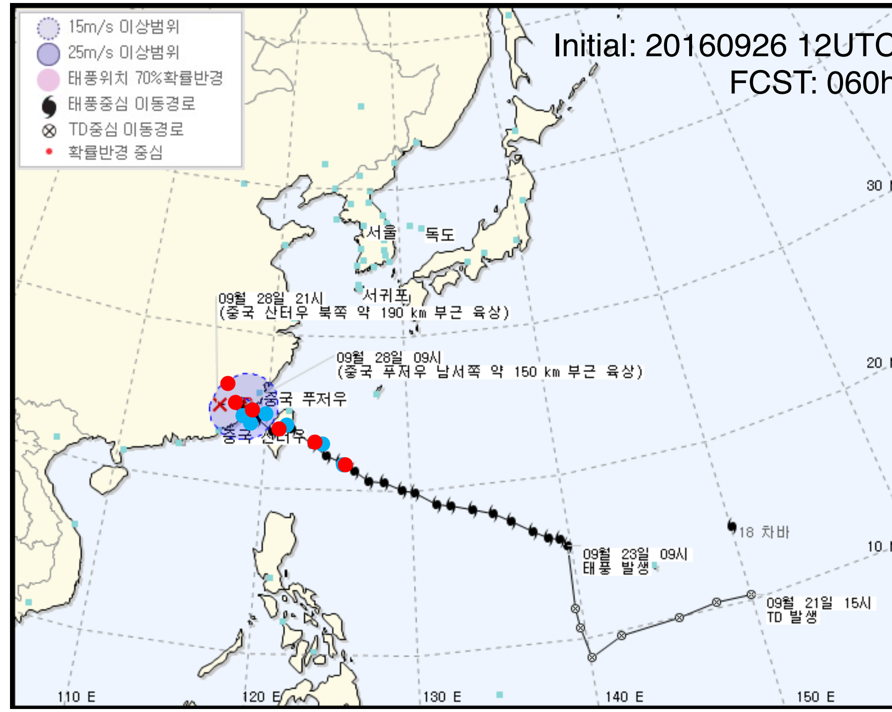

- KIM Cold (NE 120 = 25km)
- UM (N768 = 17km)

Init: 20160928 00UTC
Valid: 20161002 00/12 UTC (+96, 108h)

Forecast Comparison

KIM L50

Init : 20160928 0000UTC
Valid : 20161002 0000UTC

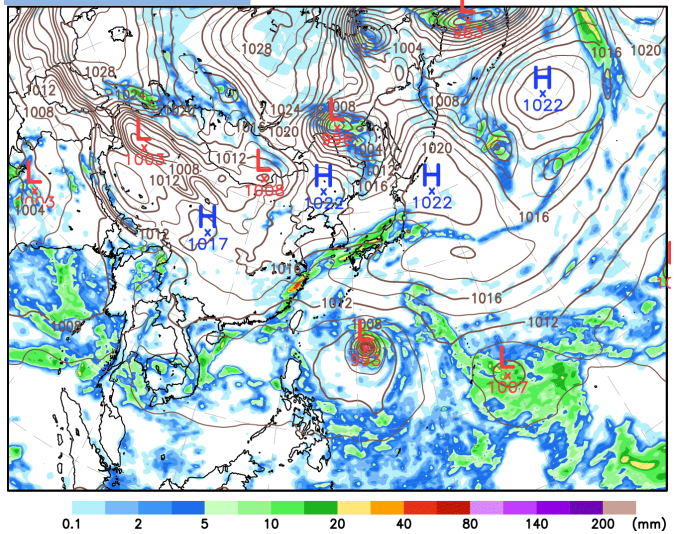

UM

Init : 20160928 0000UTC
Valid : 20161002 0000UTC

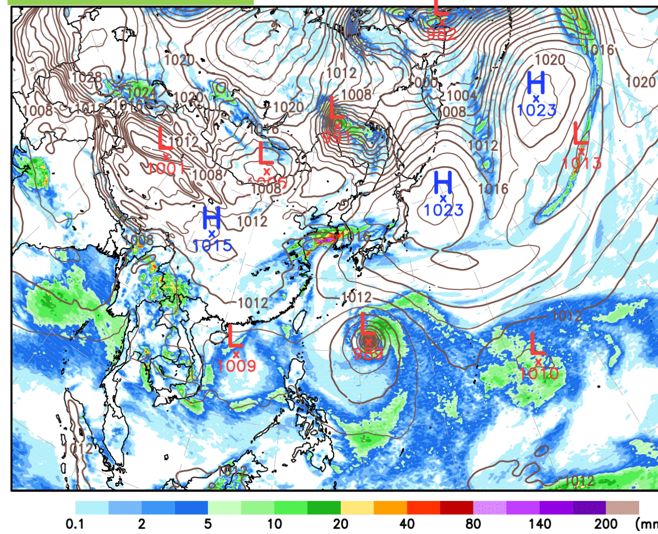

KIM 2.4 COLD ne120 L50
Surface

Init : 20160928 0000UTC
Valid : 20161002 1200UTC

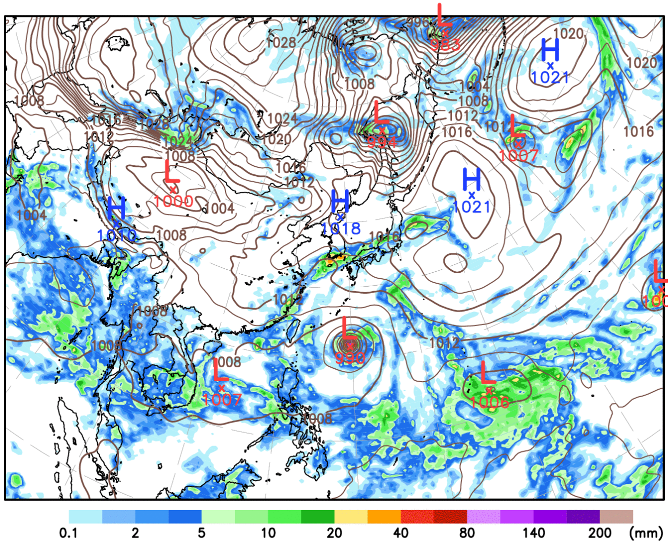

Solid line : Sea Level Pressure (hPa)
Shaded : 6 hr Accumulated precipitation (mm)

UM GDAPS N768 L70
Surface

Init : 20160928 0000UTC
Valid : 20161002 1200UTC

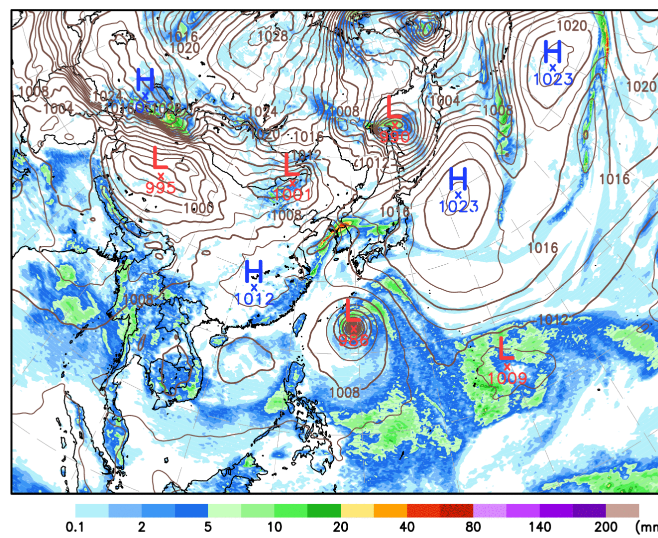

Solid line : Sea Level Pressure (hPa)
Shaded : 6 hr Accumulated precipitation (mm)

Specification of Uri System

- We have been using Uri system (Cray supercomputer) for developing KIM model.
- The system configuration is as follows.

System Configuration	Description
Inter-Connection Network	Cray Aries with Dragonfly topology
Number of Nodes	448
Processor Type	Intel 12-core Haswell
Peak Performance Per Node	998.4 GFlop/s
Total Peak Performance	447.3 TFlop/s
Memory Per Node	128 GB
Total Application Memory	56.0 TB

OpenMP parallelization, vectorization, and Hyperthreading Test at Uri System

- Experimental result
  - In order to improve the computational performance for KIM, OpenMP parallelization and vectorization (simd directives) is applied to the dynamics and physics modules.
    ➔ performance improvement ratio: about 5 ~ 10%
  - The use of both the hyperthreading and OpenMP features at Uri system can improve the computational performance of the dynamics and physics modules.
    ➔ performance improvement ratio: about 5% ~ 20%

Optimization Options	Switches	Procedure	Normal				Parallelization				Parallelization & Vectorization			
			No Hyperthreading		Hyperthreading		No Hyperthreading		Hyperthreading		No Hyperthreading		Hyperthreading	
			P720 (s)	Imprv.(%)	P1440 (s)	Imprv.(%)	P720 (s)	Imprv.(%)	P1440 (s)	Imprv.(%)	P720 (s)	Imprv.(%)	P1440 (s)	Imprv.(%)
General Optimization	-O2	RunCore_SW	0.759153		0.722203	4.87%	0.718576	5.35%	0.677127	10.80%	0.700492	7.73%	0.660644	12.98%
		rad_sw_rrtmg_driver	1.576		1.424	9.64%	1.554	1.40%	1.335	15.29%	1.4335	9.04%	1.28	18.78%
		GRIMSphysics	0.123864		0.123119	0.60%	0.118729	4.15%	0.116542	5.91%	0.125585	-1.39%	0.115856	6.47%
Processor-specific	-axCORE-AVX2 -mtune=core-avx2 -march=core-avx2	RunCore_SW	0.744678	1.91%	0.718	5.42%	0.703619	7.32%	0.660568	12.99%	0.687178	9.48%	0.654119	13.84%
		rad_sw_rrtmg_driver	1.566	0.63%	1.41	10.53%	1.555	1.33%	1.3545	14.05%	1.4225	9.74%	1.2475	20.84%
		GRIMSphysics	0.122949	0.74%	0.123763	0.08%	0.119161	3.80%	0.117364	5.25%	0.125831	-1.59%	0.114636	7.45%
Inter-Procedural Optimization (IPO)	-axCORE-AVX2 -mtune=core-avx2 -march=core-avx2 -ipo	RunCore_SW	0.741051	2.38%	0.707695	6.78%	0.69411	8.57%	0.658949	13.20%	0.681339	10.25%	0.643127	15.28%
		rad_sw_rrtmg_driver	1.622	-2.92%	1.388	11.93%	1.5595	1.05%	1.3885	11.90%	1.434	9.01%	1.2825	18.62%
		GRIMSphysics	0.121153	2.19%	0.121492	1.92%	0.117492	5.14%	0.115932	6.40%	0.125483	-1.31%	0.113831	8.10%

MPI vs. OpenMP Scaling Analysis

- Experimental result
  - n-Way means the number of MPI processes times the number of OpenMP threads
  - Computational performance according to the number of threads
    ➔ As the number of MPI processes increases, when two threads is used, the computational performance of KIM is optimal regardless of the use of hyperthreading.

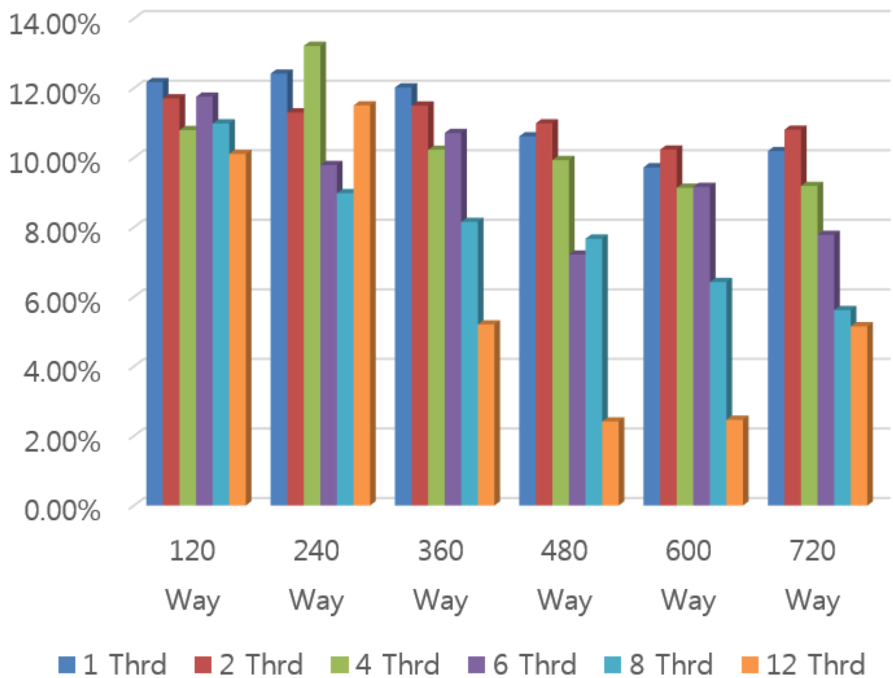

Without Hyperthreading

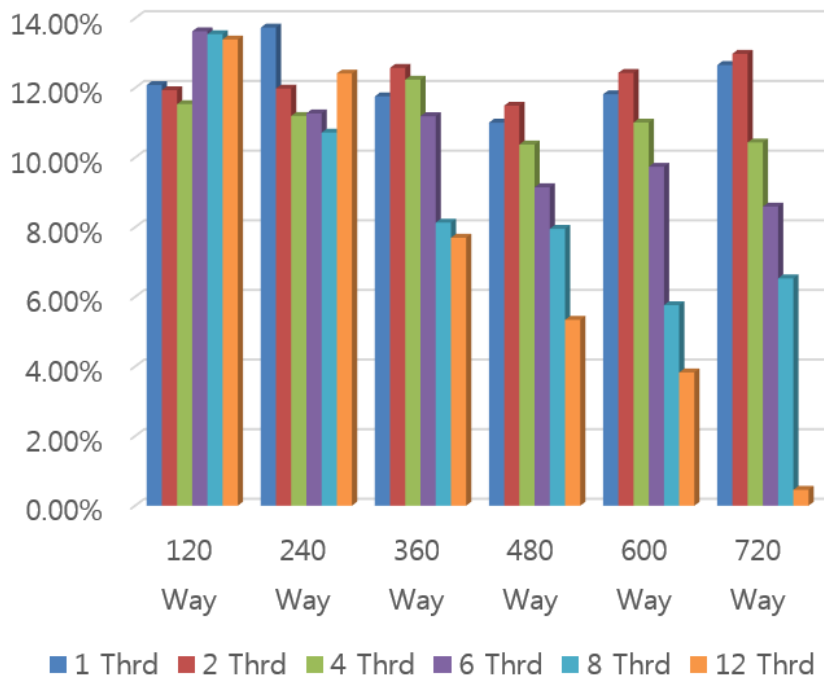

With Hyperthreading

Hybrid MPI/OpenMP Program

- This figure is the extension to the scaling analysis using the processors more than 720-way, including 2400, 3600, and 5400-way.
- This figure also represents that it is easy to determine the optimal number of MPI tasks and OpenMP threads when the network condition for the MPI communication is not congested.
  - The use of two OpenMP threads provides the optimal performance of KIM model.

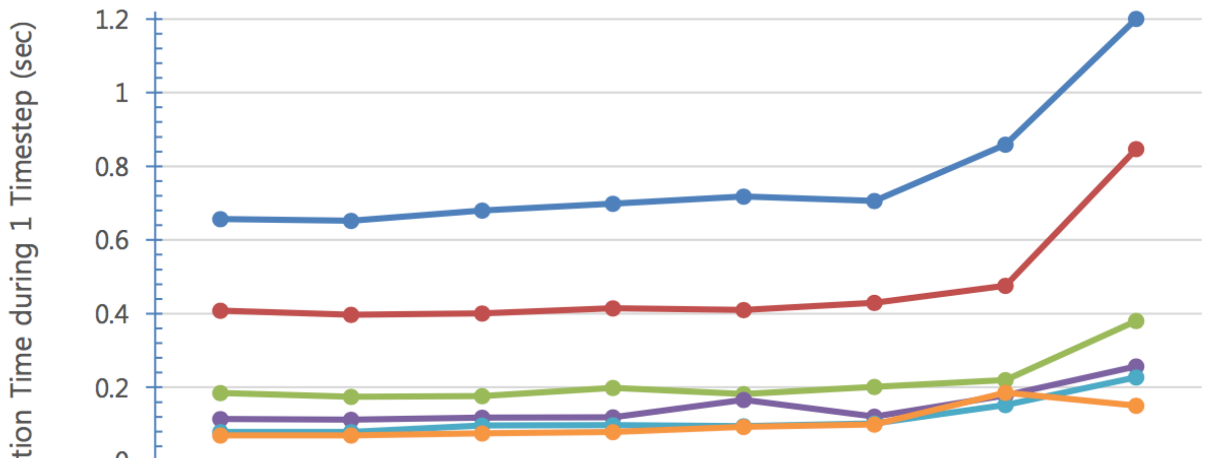

	T1	T2	T4	T6	T8	T12	T24	T48
P720	0.6565	0.6525	0.68	0.699	0.718	0.7065	0.859	1.2005
P1200	0.4085	0.3975	0.4005	0.415	0.4105	0.43	0.476	0.847
P2400	0.185	0.1745	0.1765	0.199	0.1815	0.201	0.22	0.3805
P3600	0.1145	0.112	0.118	0.119	0.1665	0.1205	0.1785	0.257
P4800	0.079	0.079	0.0965	0.097	0.0945	0.102	0.152	0.2275
P5400	0.07	0.0695	0.075	0.079	0.0925	0.0995	0.1855	0.15

Hybrid MPI/OpenMP Program

- It is difficult to determine the optimal number of MPI tasks and OpenMP threads when the network condition for the MPI communication is poor.
- In case of 3600, 4800, and 5400-way, the use of 48 OpenMP threads can offer the optimal performance (i.e., it can provide better performance compared to that of 2 OpenMP threads.)

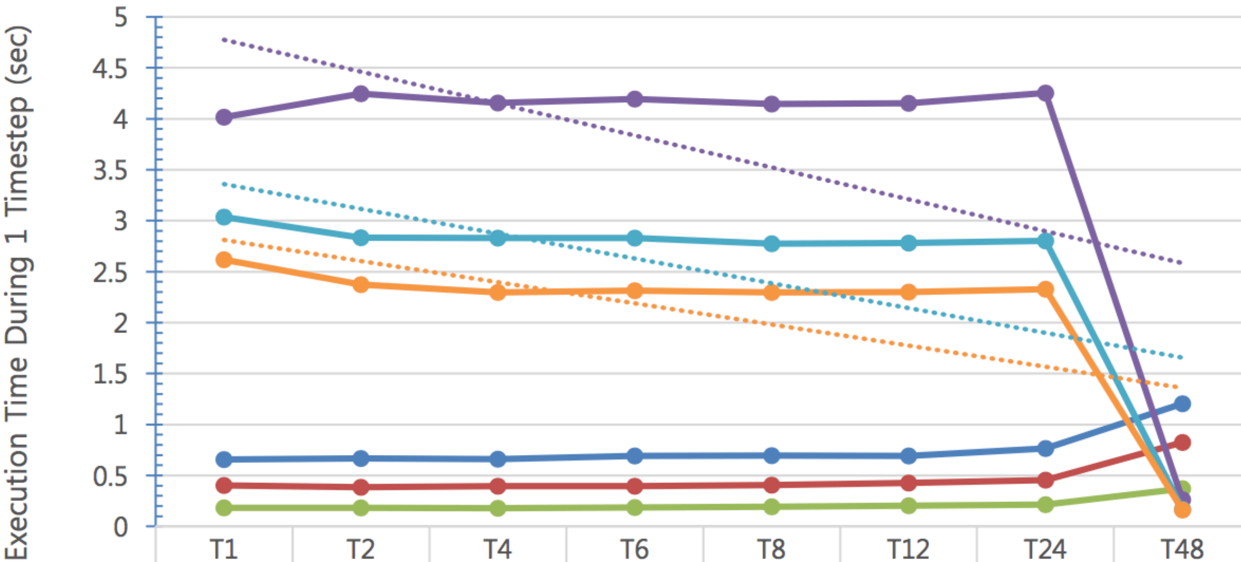

	T1	T2	T4	T6	T8	T12	T24	T48	
P720	0.658432	0.666229	0.65978	0.691924	0.696703	0.693653	0.764831	1.205229	
P1200	0.402568	0.386661	0.394907	0.39622	0.404949	0.425839	0.453373	0.825246	
P2400	0.183305	0.182966	0.18078	0.1855	0.193864	0.203246	0.21272	0.372271	
P3600	4.015764	4.24522	4.154712	4.194661	4.14572	4.153593	4.252746	0.262881	→ 15.28x speedup
P4800	3.03661	2.832653	2.828737	2.828653	2.775508	2.782144	2.803932	0.167602	→ 18.11x speedup
P5400	2.617898	2.373492	2.297975	2.312941	2.296847	2.301282	2.328661	0.163127	→ 16.04x speedup

Number of Thread per MPI Tasks

Conclusion and Future Works

- We can achieve optimal performance when using both hyperthreading and OpenMP features (parallelization and vectorization), compared to employing only the OpenMP features.

- OpenMP enabled KIM with more threads performs significantly better than that with fewer threads in congested network conditions.

- We plan to apply OpenMP offloading feature using computational accelerators, such as intel Xeon Phi MIC and GPU to KIM model.

KIAPS

KOREA INSTITUTE OF
ATMOSPHERIC PREDICTION SYSTEMS

Thank You!

www.kiaps.org