The Giga-Large Eddy Simulation-2 Experiment

Donald Dazlich[1], Ian Baker[1], Peter Blossey[2], Scott Denning[1], John Helly[3], Marat Khairoutdinov[4], Chin-Hoh Moeng[5], Hugh Morrison[5], Robert Pincus[6], David Randall[1]

[1]Department of Atmospheric Science, Colorado State University, Fort Collins, Colorado USA
[2]Department of Atmospheric Sciences, University of Washington, Seattle, Washington USA
[3]San Diego Supercomputer Center/University of California San Diego/Scripps Institution of Oceanography, San Diego, California USA
[4]School of Marine & Atmospheric Science, Stony Brook University, Stony Brook, NY USA
[5]National Center for Atmospheric Research, Boulder, Colorado USA
[6]CIRES University of Colorado/NOAA Earth System Research Lab, Boulder, Colorado USA

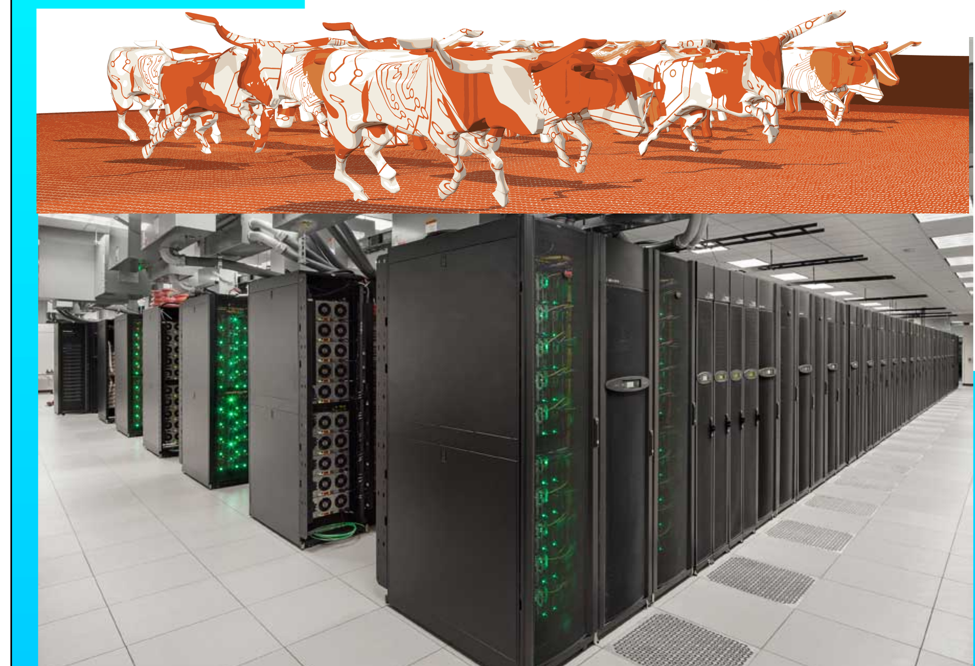
TACC STAMPEDE


NICS KRAKEN


SDSC GORDON

Introduction

Cloud processes remain a substantial uncertainty in our knowledge of the climate system. Large-eddy simulation (LES) provides the most detailed treatment of clouds available by resolving enough of the large eddies that those unresolved are satisfactorily treated with simple theory. By performing LES over a large enough domain not only are individual clouds simulated in detail, but we can simulate dynamically-organized cloud systems over an area comparable to a single grid-cell of a climate model. In our simulations, we use a cloud-resolving model (CRM) to represent a 205km x 205km domain with a horizontal grid spacing of 100m (2048x2048 grid-cells in the horizontal plane), and discretized into 256 vertical layers resulting in model variables predicted at one billion locations – hence the name GigaLES. No field campaign could observe these systems in such detail, but GigaLES simulation can serve as a complementary source of data from which we can develop our understanding of cloud systems.

The first GigaLES simulation was performed in 2009 (Khairoutdinov et al 2009). It used an idealized representation of the large-scale atmospheric environment, radiative heating rates were prescribed in a time-invariant and spatially homogeneous fashion, a basic one-moment microphysics treatment was used, and the model was run for one simulated day. They found a convergence of the statistical properties between 100m and 200m grid spacing. Now we wish to simulate deep-convective cases over both ocean and land, with more sophisticated model parameterizations including computed radiation, and forced with a time-varying atmospheric environment as measured by field campaigns and over multiple diurnal cycles and test the robustness of the previous results.

Model

SAM6.10.4 : This CRM predicts momentum using the anelastic equations of motion. The prognostic thermodynamic variable is the liquid/ice water moist static energy. The model is configured to run with:

- Two-moment treatment of cloud microphysical processes predicting water vapor, and the mass and number of five condensed water species.
- RRTMG interactive radiation.
- 1.5 order closure of the turbulence kinetic energy budget to compute sub-grid scale fluxes.
- Ultimate Macho 5 scheme to advect all scalar quantities monotonically with 5th order accuracy.

The model was modified to include:

- Simple Biosphere 3 (SiB3), a prognostic land-surface model with photosynthetic control of fluxes to the atmosphere.
- New cloud optical properties based on CAM5, including radiatively-active snow and using predicted size distributions from the two-moment microphysics.
- Lagrangian Parcel Tracker (LPT) diagnostic package that predicts the trajectories of user-defined parcels.

The CRM is configured to run with doubly-periodic boundary conditions, e.g. anything carried out the east boundary of the domain reenters at the west boundary. Thus, real geographic domains must be somewhat idealized.

Simulations

Marine Simulation – TWP-ICE

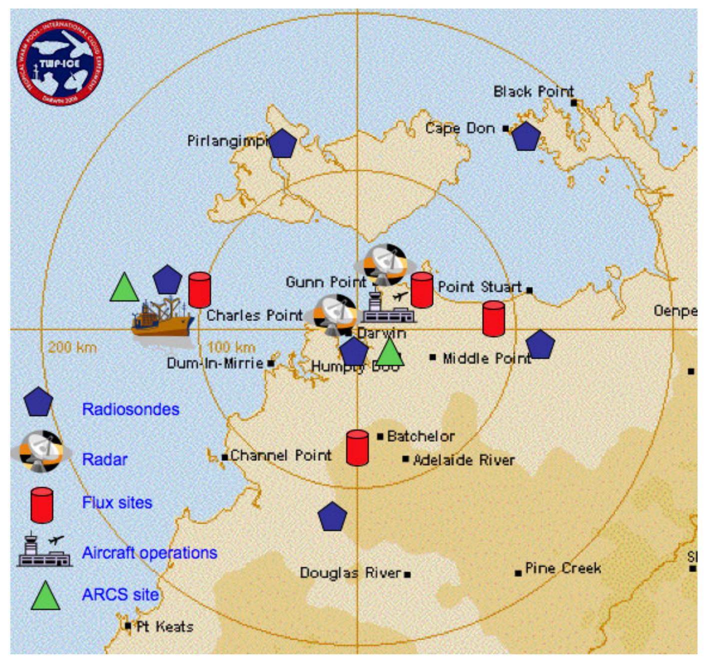

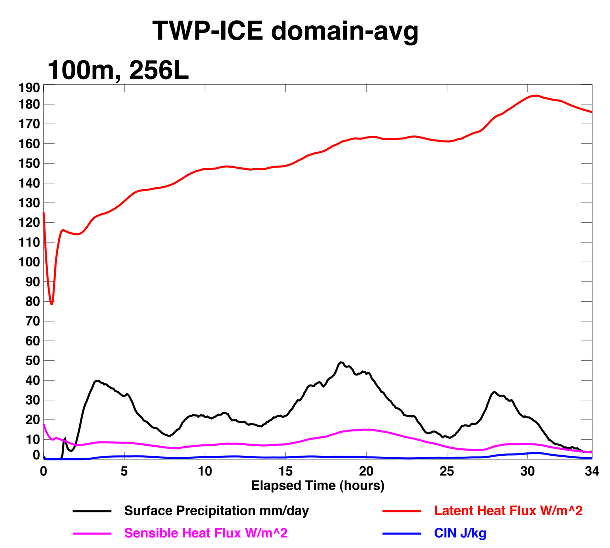

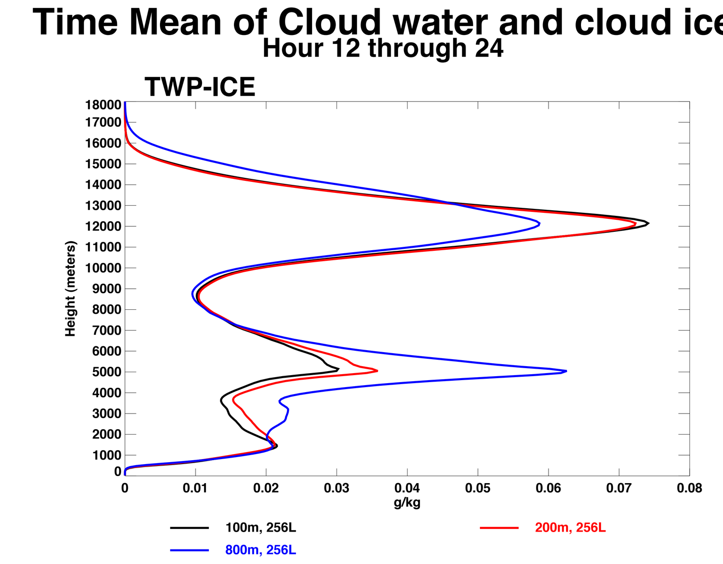

Tropical Warm Pool – International Cloud Experiment (TWP-ICE) – simulations in progress.

- Location – 130.9E,12.4S
- Period – 18 Jan 2006 – 4 Feb 2006; first few days for 100m grid spacing.
- Features enhanced and suppressed convective periods.
- The model is run according to the CRM intercomparison specifications in Fridland et al, 2010.

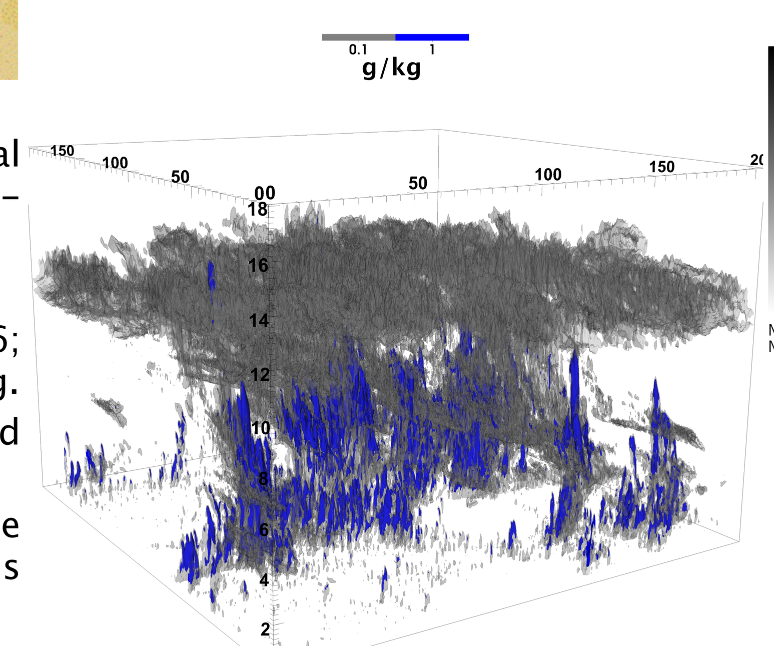

Cloud-condensate field at hour 24

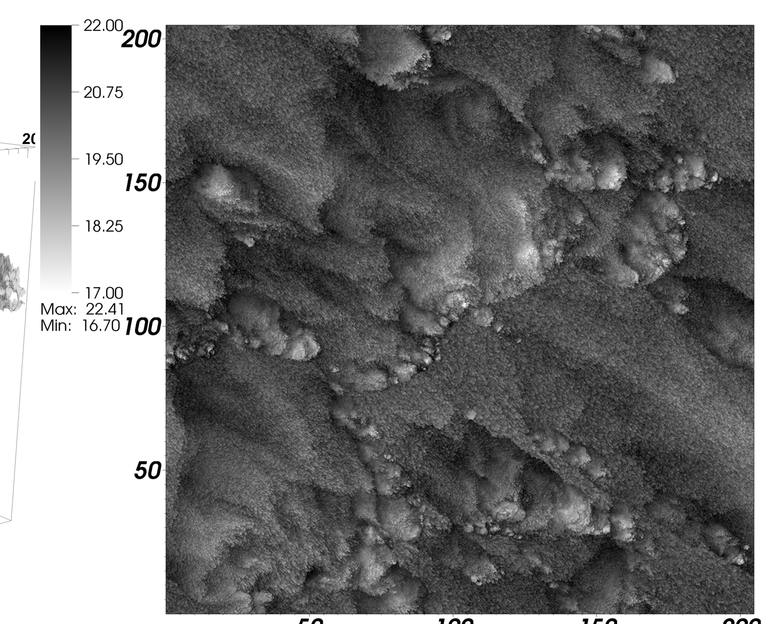
25m water vapor mixing ratio at hour 18.5, 100m

Land Simulation – MC3E

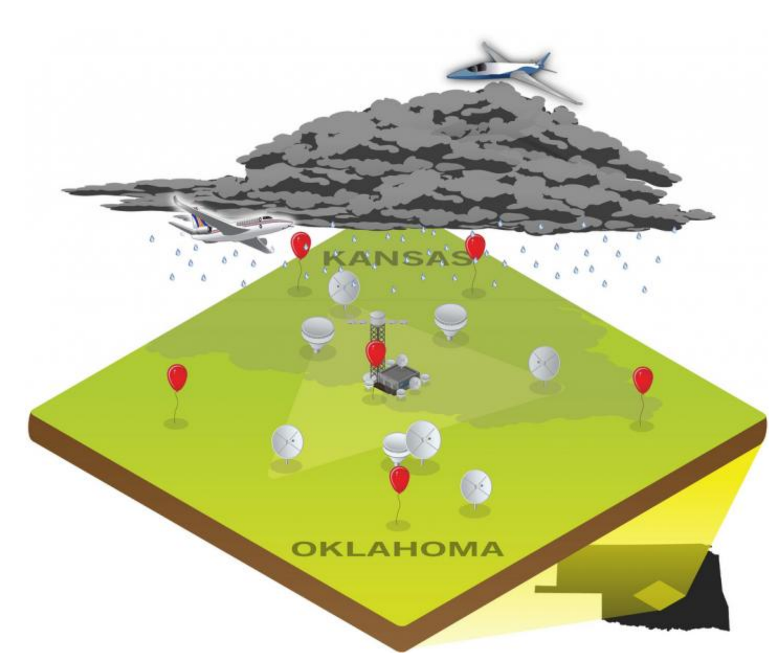

Midlatitude Continental Convective Clouds Experiment (MC3E) – simulations to be done.

- Location – 97.5W, 36.5N
- Period – 22 Apr 2011 – 7 Jun 2011; 23 May – 26 May for 100m grid spacing.
- The CRM is forced with advective tendencies of temperature and water vapor derived from the field campaign observations.

The addition of SiB3 to the model permits a simulation where a number of land variables are predicted and feedback to the fluxes passed to the CRM. This increases the degrees of freedom of the model.

A key features of GigaLES-2 Land is the heterogenouus land surface. On the left is the actual distribution (green=grassland, yellow=agriculture); on the right is the random distribution assigned in the model.

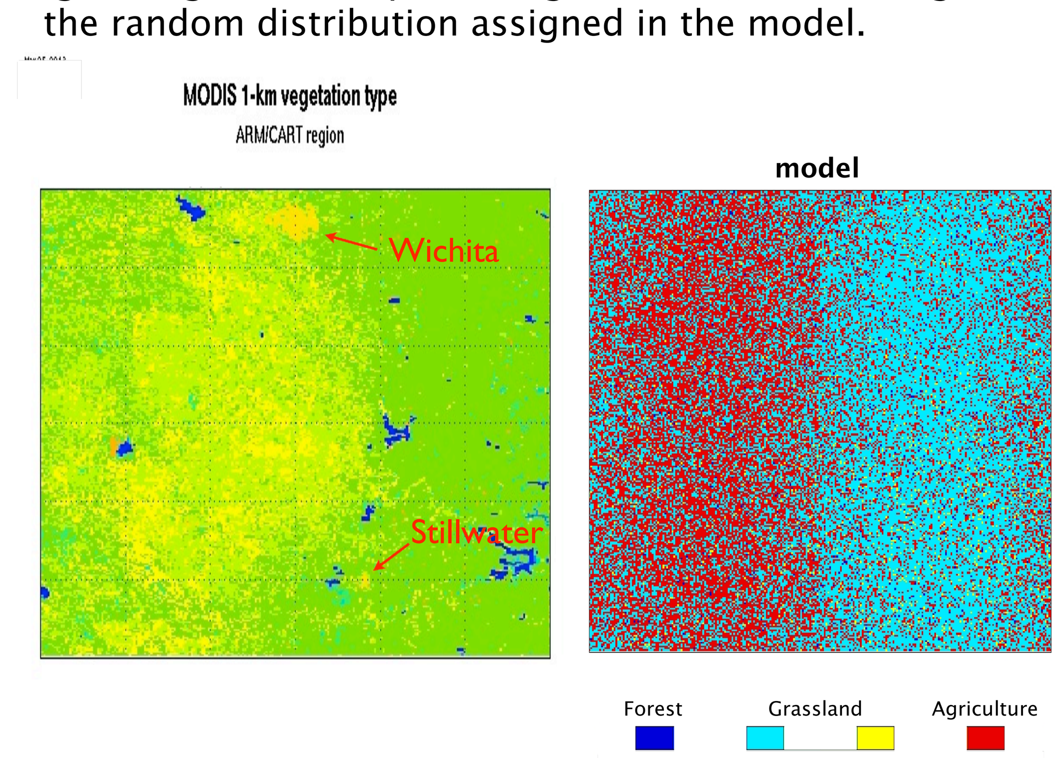

Computational Aspects

The GigaLES-2 configuration is a 2048x2048 gridcell domain in the horizontal with a grid spacing of 100m, and 256 layers in the vertical with grid spacing varying from 50m near the surface, to 100m near the tropopause, to 300m at the model top of 27km. The fine grid spacing requires that the predicted variables be updated with a timestep of 2 seconds. The computation and memory required by this configuration requires that the domain be distributed across 1024 processors of a supercomputer. MPI message-passing is used to share data between processors.

Numerical methods that perform adequately on a small number of computational processors may become bottlenecks on very large number of processors needed to perform a GigaLES simulation. Message passing and I/O are particularly susceptible to poor scaling. Two parts of SAM that scaled poorly were the elliptic solver and the 3D array write. In the former a domain decomposition transpose that cost O(n) messages per process, where n is the number of processes, was replaced with a scheme that cost $O(n^{1/2})$, and the Fourier transform package was replaced with FFTW, faster and eliminating redundent computations. In the latter, the serial (one process) write of the global 3D array was replaced with a distributed write using parallel NetCDF-4. This also eliminated the need for a conversion to NetCDF format in the post-processing.

SCALING – grid cells per process constant

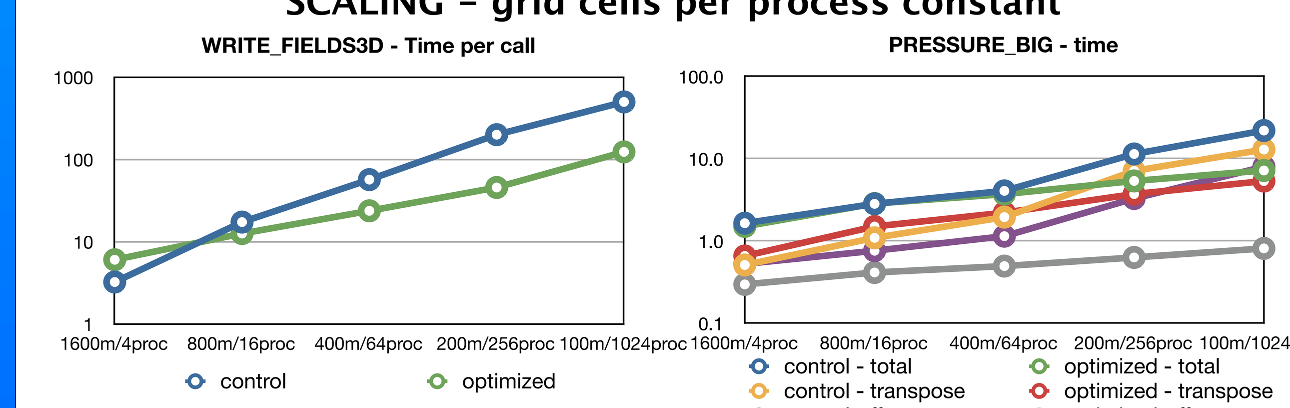

Acknowledgements

This study was supported by the National Science Foundation's Science and Technology Center for Multi-scale Modeling of Atmospheric Processes (CMMAP), managed by Colorado State University under cooperative agreement No. ATM-0425247.

The GigaLES-2 computations used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575, including the Cray XT5 'Kraken' at the National Institute of Computational Sciences, the Linux cluster 'Gordon' at the San Diego Supercomputer Center, and the Dell Linux Cluster 'Stampede' at the Texas Advanced Computing Center.

References

Fridlind, A. M., A. S. Ackerman, J. Petch, P. R. Field, A. Hill, G. G. McFarquhar, S. Xie, and M. Zhang (2010), ARM/GCSS/SPARC TWP-ICE CRM intercomparison study, NASA Tech. Memo., NASA TM-2010-215858, 24 pp. http://pubs.giss.nasa.gov/docs/2010/2010_Fridlind_etal.pdf

Khairoutdinov, M. F., and D.A. Randall, 2003: Cloud-resolving modeling of the ARM summer 1997 IOP: Model formulation, results, uncertainties and sensitivities. J. Atmos. Sci., 60, 607-625.

Khairoutdinov M. F., S. K. Krueger, C.-H. Moeng, P. A. Bogenschutz, and D. A Randall, 2009: Large-eddy simulation of maritime deep tropical convection, J. Adv. Model. Earth Syst., Vol. 1, Art. #15, 13 pp., doi:10.3894/JAMES.2009.1.15