Adjoint Based Scientific Applications for Large Scale Atmospheric Models and their Performance on General Purpose GPUs

presented

by

Kumaresh Singh

PhD Candidate

Department of Computer Science

Outline

- Motivation
- Challenges
- Introduction
- GEOS-Chem Adjoint Model
- Chemistry using CUDA
- Conclusion
- Future Work

Motivation

- Adjoint models are powerful tools widely used in meteorology and oceanography for applications
  - data assimilation
  - model tuning
  - sensitivity analysis
  - determination of singular vectors
- Ironical - Working on improving air quality by running model codes on power hungry architectures

Challenges

- High learning curve to acquaint with large scale air quality models
- Need to understand working of each single subroutine and parameters involved
- Rigorous work of constructing and testing each science process adjoint individually and integrating those into a consistent model
- Current GPGPUs lack some conventional shared memory features and are less explored

Introduction

- Sensitivity Analysis

  Adjoint model is efficient in calculating sensitivities of a few output variable or metrics with respect to a large number of (input) parameters

- Data Assimilation

  4D-Var data assimilation allows the optimal combination of three sources of information: an a priori (background), estimate of the state of the atmosphere, knowledge about the physical and chemical processes

Data Assimilation Framework

- At iteration 0, $x_0 = c_p^0$
- At each subsequent iteration k ($k \geq 1$),

  $x_{k+1} \leftarrow$ L-BFGS (x_k, f, g)

  $c_{op}^0 \leftarrow x_{k+1}$

  (f, g) $\leftarrow$ reverse_mode (c_{op}^0, Observation_Chk)

  where, f is the cost function and g is the gradient of the cost function.

- In our test case, the cost function and its gradient are defined as:

  $f = (1/2)\sum(c_{op}^{k,m} - c_0^{k,m})^T R\text{--}_k^{-1} (c_{op}^{k,m} - c_0^{k,m}) + (1/2)\sum(c_{op} - c_b)^T B^{-1} (c_{op} - c_b)$

  $g = \sum R\text{--}_k^{-1} (c_{op}^{k,m} - c_0^{k,m}) + \sum B^{-1} (c_{op} - c_b)$

Time Distribution of GC processes

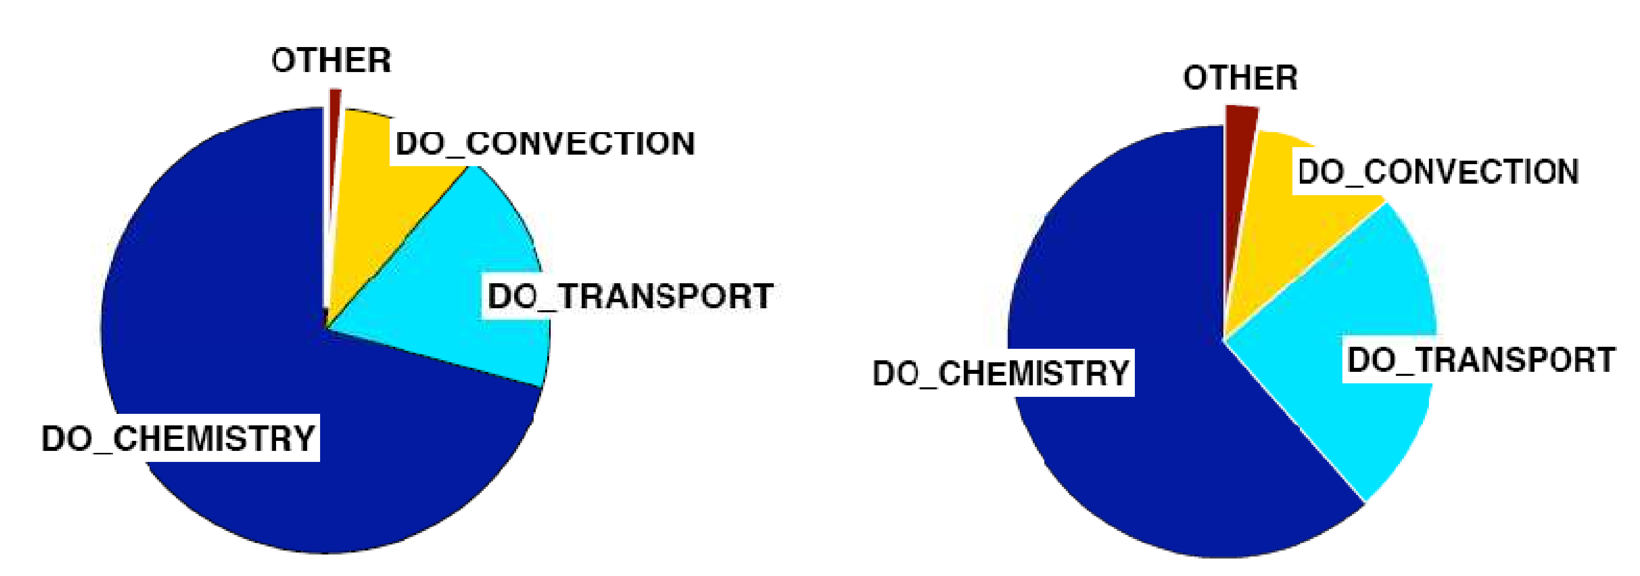

(a) 1 cpu (12 min 46 sec) (b) 8 cpus (2 min 7 sec)

GEOS-Chem simulation times on one and on eight cores. The majority of the computational time is spent in three processes: chemistry, transport, and convection.

GEOS-Chem

- A global 3-D model of atmospheric composition driven by assimilated meteorological observations from the Goddard Earth Observing System (GEOS)
- Used by research groups worldwide for
  - assessing intercontinental transport of pollution
  - evaluating consequences of regulations and climate change on air quality
  - comparison of model estimates to satellite observations and field measurements
  - fundamental investigations of tropospheric chemistry

GEOS-Chem Users

The International GEOS–Chem User Community (N. America)

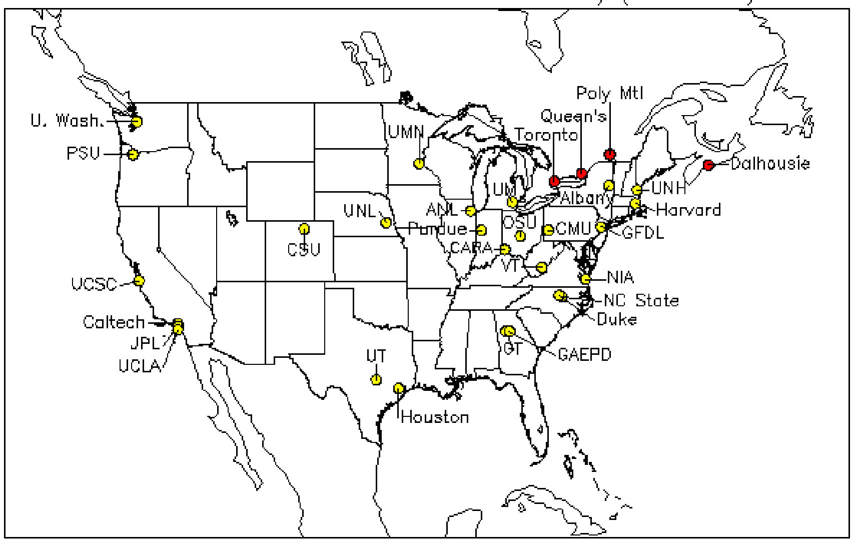

The International GEOS–Chem User Community (Europe)

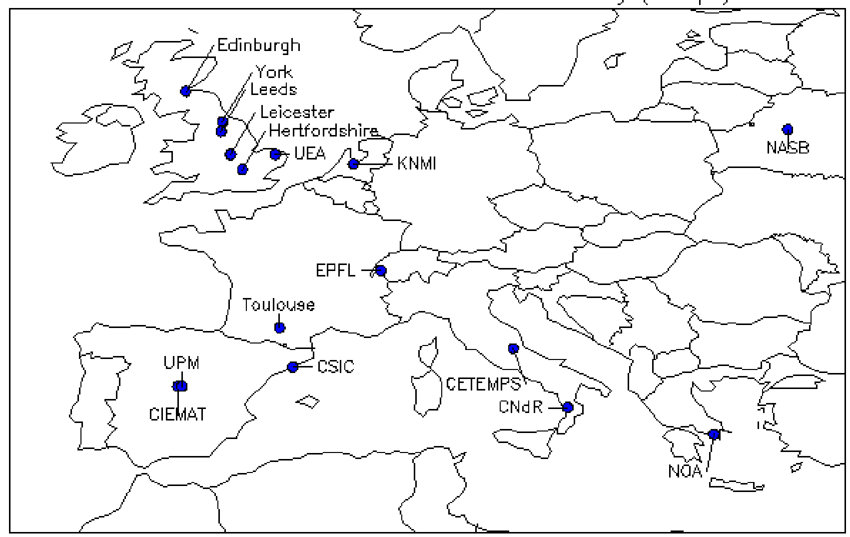

The International GEOS–Chem User Community (Asia & Oceania)

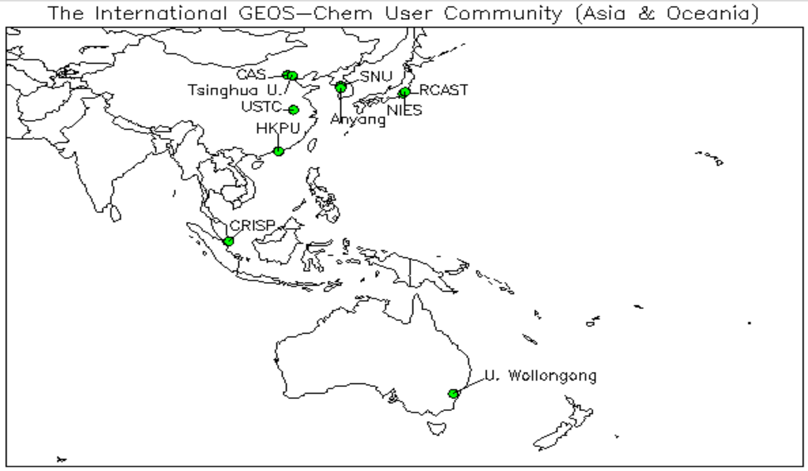

GEOS-Chem Adjoint (GCv7_ADJ)

- Created an adjoint model of geos-4 v7 of GEOS-Chem
  - Tested each science process adjoint separately
  - Consistency check after integrating all processes together
  - Completely parallelized adjoint code
- Added 4-D variational data assimilation and sensitivity analysis capabilities
- Provided with choices of operations to choose from as per the need, plug-n-play functions for cost function calculations
- Adjoint code quite similar to forward mode – same coding convention

Forward and Adjoint Code Flow

Forward Mode	Adjoint Mode
CONVERT_UNITS(kg->vv)	**CALL DO_WETDEP_ADJ**
CALL DO_UPBDFLX	Read_CHEM_CHK(Date,Time) **CALL DO_CHEMISTRY_ADJ** ▪ *Emission, dry deposition handled inside chemistry.*
CALL DO_TRANSPORT	CONVERT_UNITS(vv->kg)
CALL DO_PBL_MIX	Read_CONV_CHK(Date,Time) **CALL DO_CONVECTION_ADJ**
Make_CONV_CHK(Date,Time) **CALL DO_CONVECTION**	**CALL DO_PBL_MIX_ADJ**
CONVERT_UNITS(vv->kg)	**CALL DO_TRANSPORT_ADJ**
CALL DO_DRYDEP CALL DO_EMISSIONS ▪ *Updating emission and dry deposition rates.*	**CALL DO_UPBDFLX_ADJ**
Make_CHEM_CHK(Date,Time) **CALL DO_CHEMISTRY**	CONVERT_UNITS(kg->vv)
CALL DO_WETDEP	

Validation results

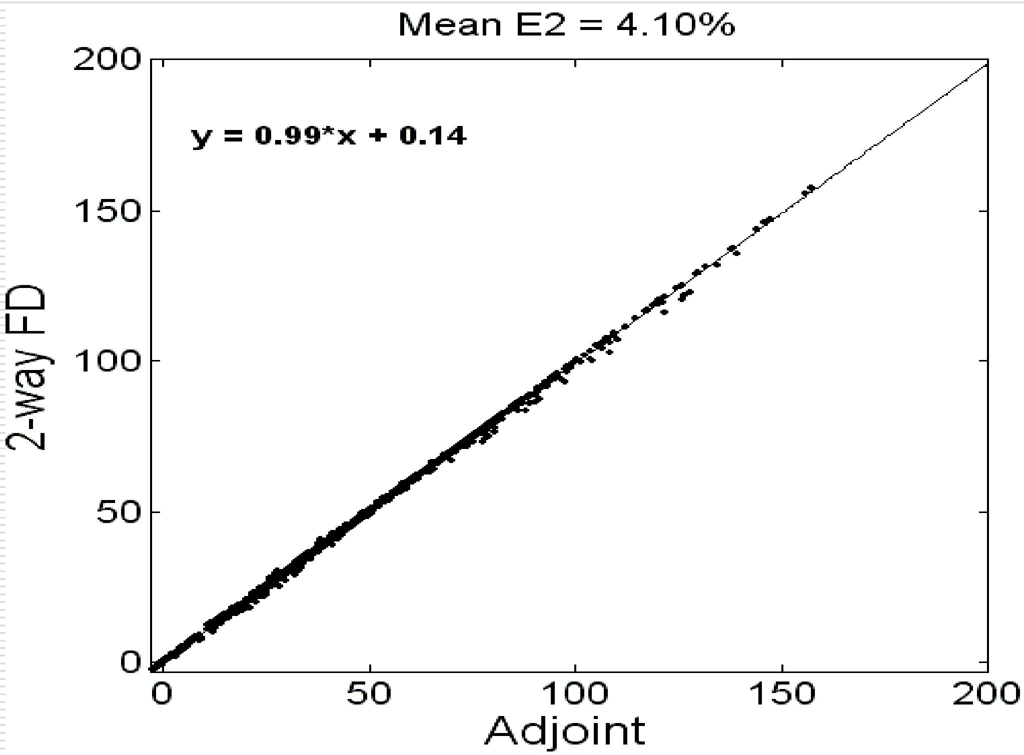

Scattered plot of adjoint vs. central finite difference values over all the grid points, generated by running GEOS-Chem v7 adjoint, **chemistry** only simulation for 6 days from 2001/04/01 to 2001/04/07, for SO4 with respect to NOx concentrations, layer 10

Validation results

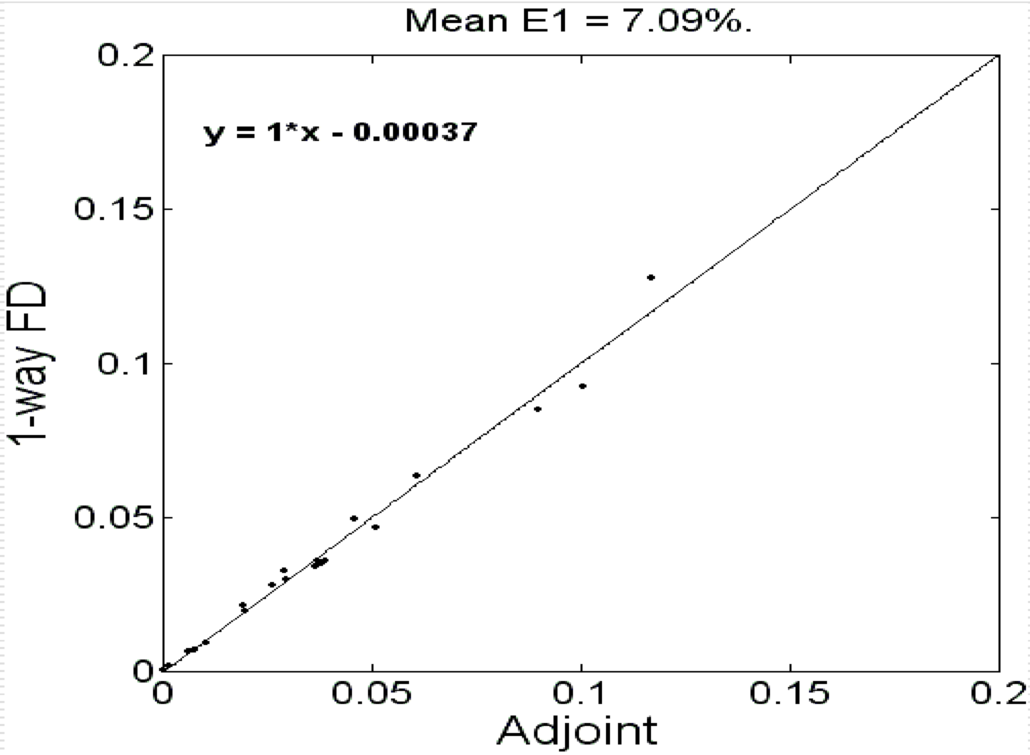

Scattered plot of 1-way finite difference vs. adjoint values generated by running GEOS-Chem v7 adjoint, **advection** only, for 2 days from 2001/07/01:00 to 2001/07/03:00, for NOx concentrations (continuous adjoint).

Validation results

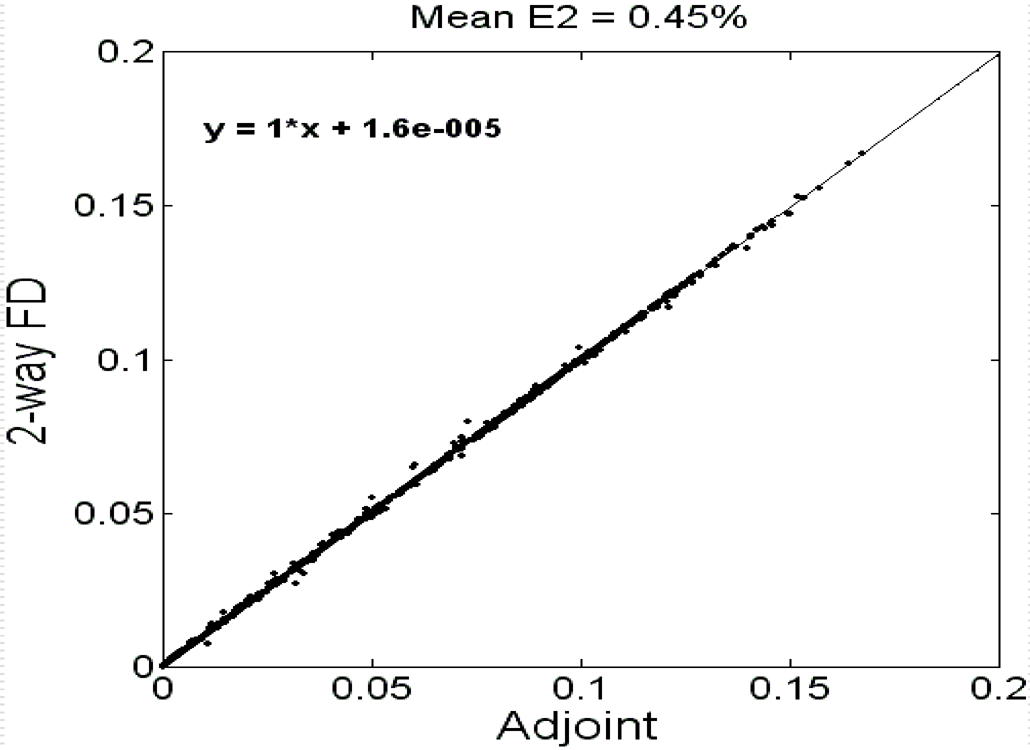

Scattered plot of 2-way finite difference vs. adjoint values generated by running GEOS-Chem v7 adjoint, **convection** only, for 6 days from 2001/07/01 to 2001/07/07, for NOx concentrations. A perturbation was introduced at layer 2 and was tracked at layer 9

Validation results

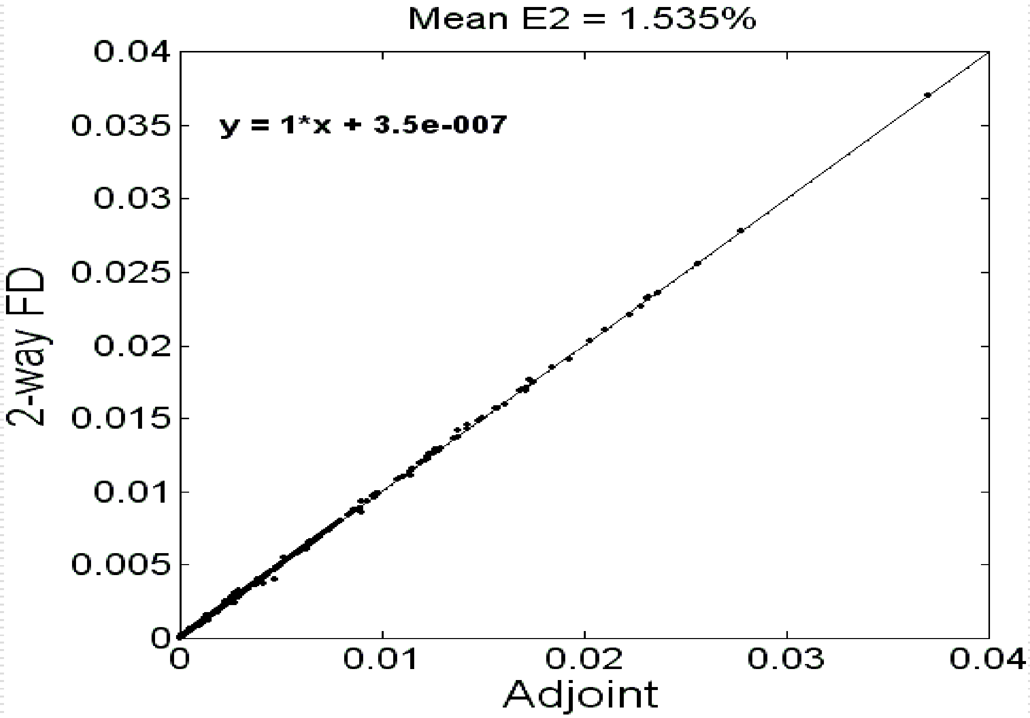

Scattered plot of central finite difference vs. adjoint values generated by running GEOS-Chem v7 adjoint, **wet deposition** only, for one week from 2001/07/01 to 2001/07/08. We consider the wet deposition process acting on H2O2 concentrations. The perturbation was introduced at layer H=9 and measured at layer L=5

Validation results

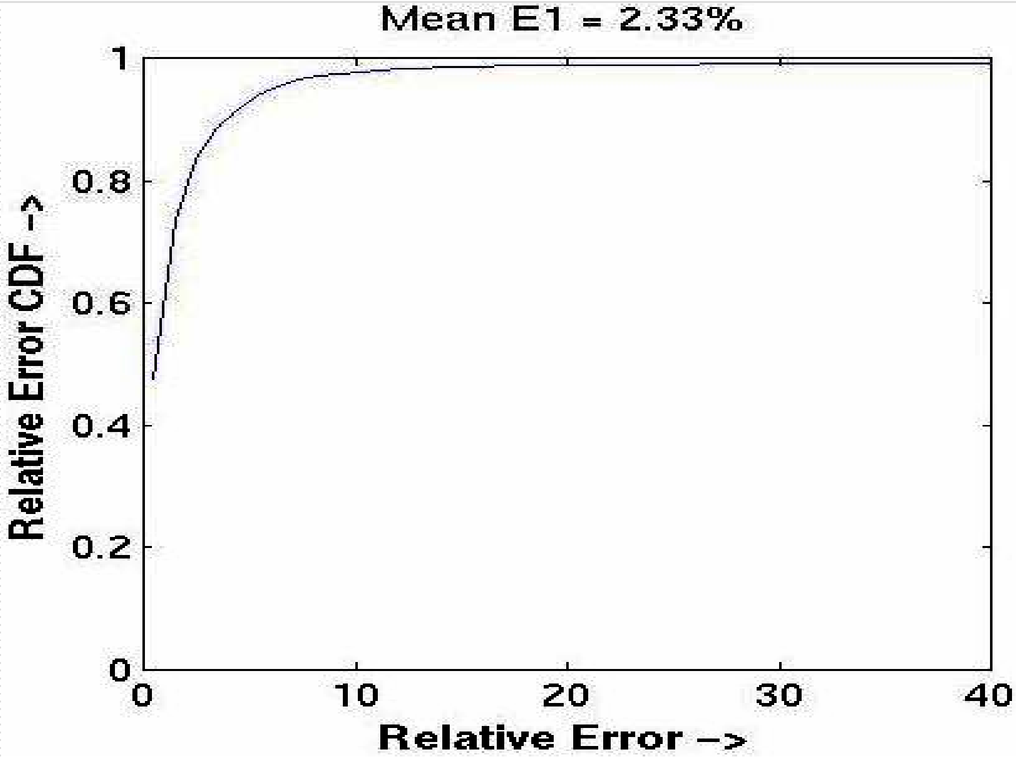

1-way finite difference vs. adjoint relative error cumulative distribution function plot generated by running GEOS-Chem v7 adjoint **emissions/dry-deposition** only, 2 days from 2001/07/01:000000 to 2001/07/03:000000, for changes in Ox concentrations with respect to NOx emissions

Adjoint Model Speed-up Graphs

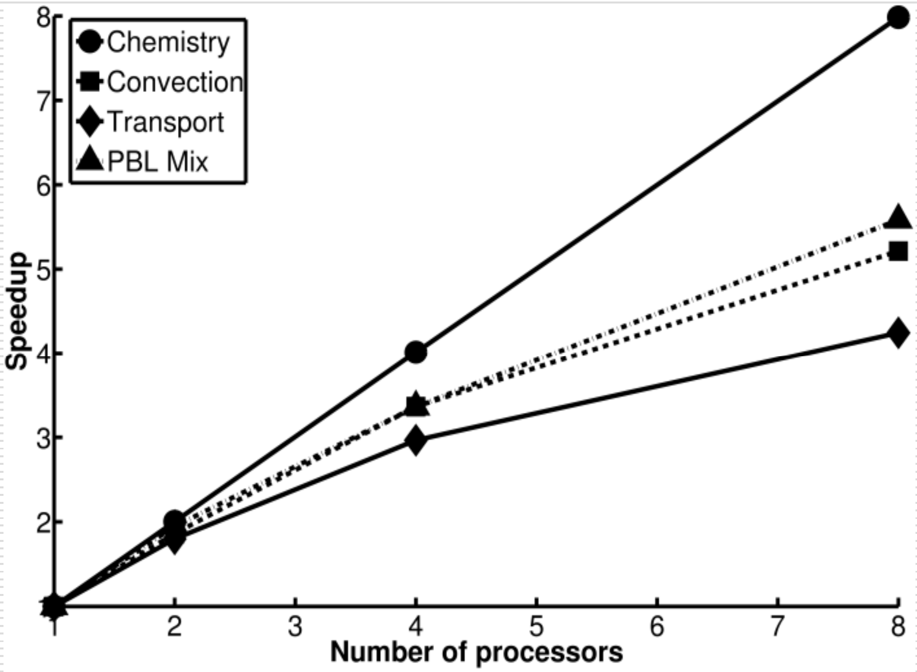

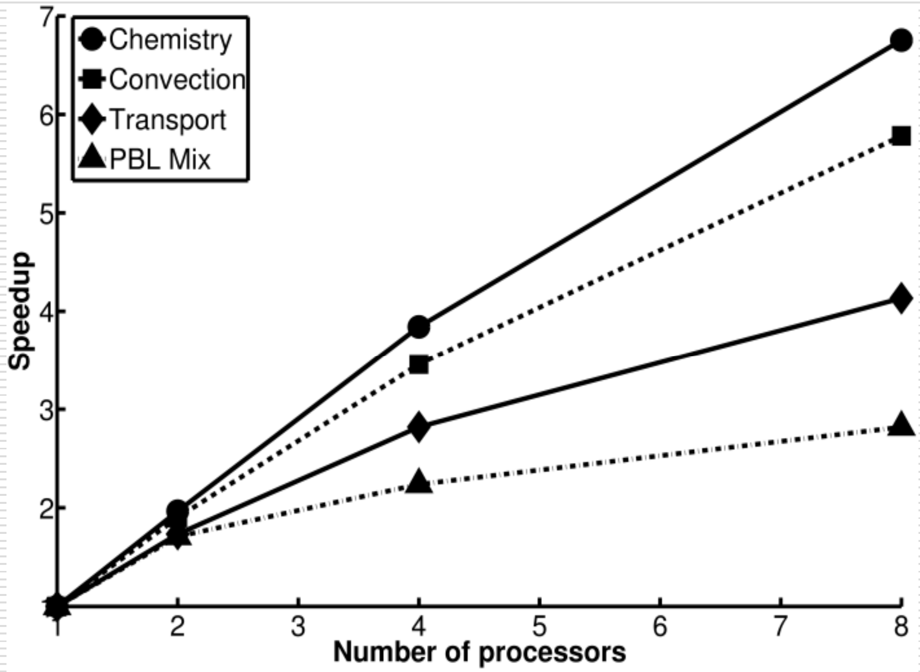

Speedup graphs for chemistry, convection, advection and planetary boundary mixing subroutines in forward(*left*) and adjoint(*right*) mode on 1, 2, 4 and 8 processors. The simulation window for this analysis was 24 hours performed on July 2001 GEOS-Chem data

Sensitivity Analysis (emission species)

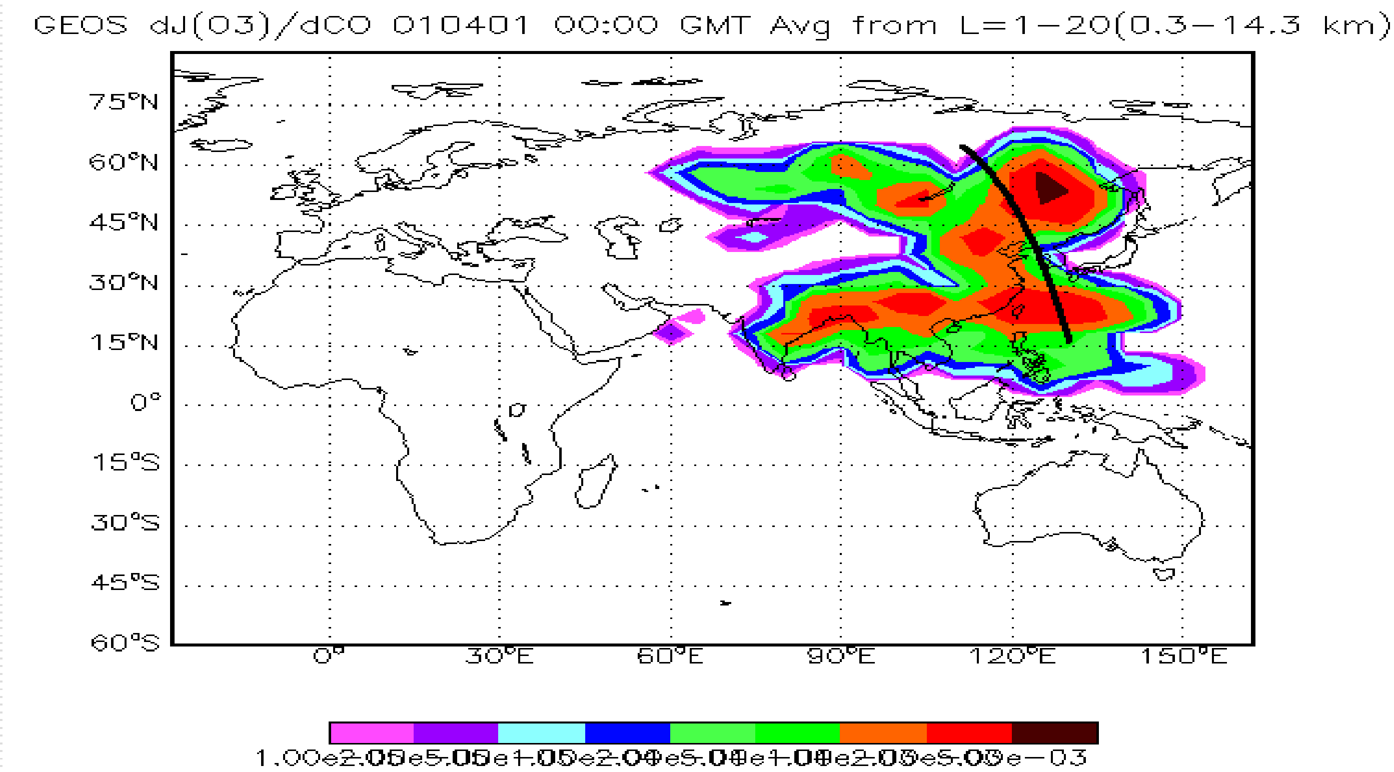

Sensitivity of the O3 column measured by TES with respect to the CO over Asia on April 1st, 2001

Sensitivity Analysis (emission species)

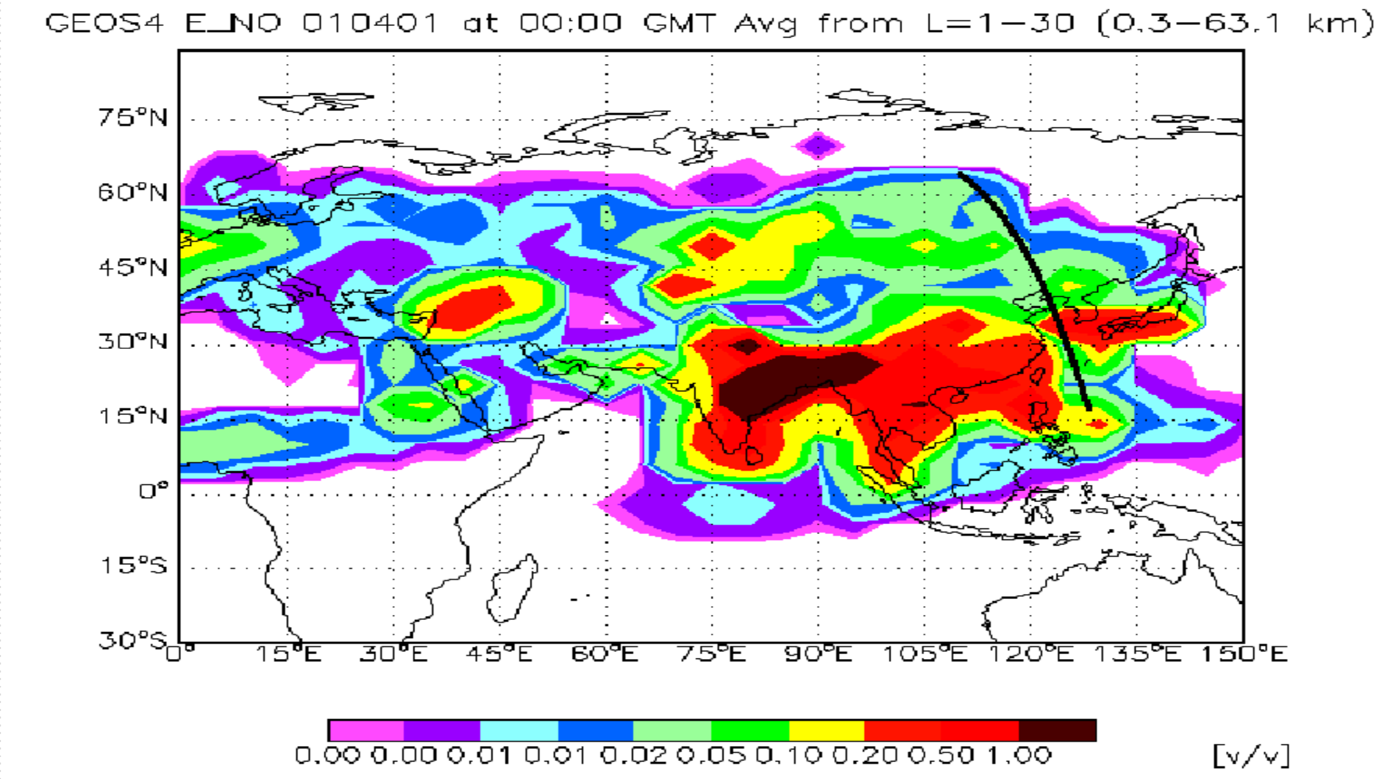

Sensitivity of the O3 column measured by TES with respect to total NOx emissions over Asia on April 1st, 2001

Sensitivity Analysis (tracer species)

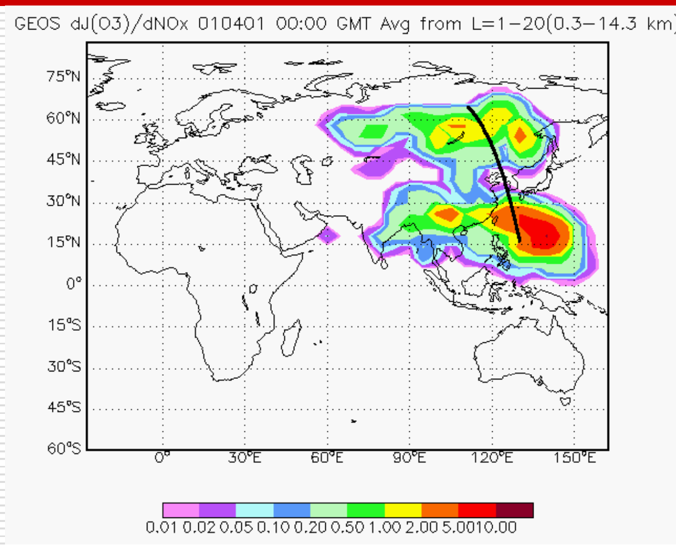

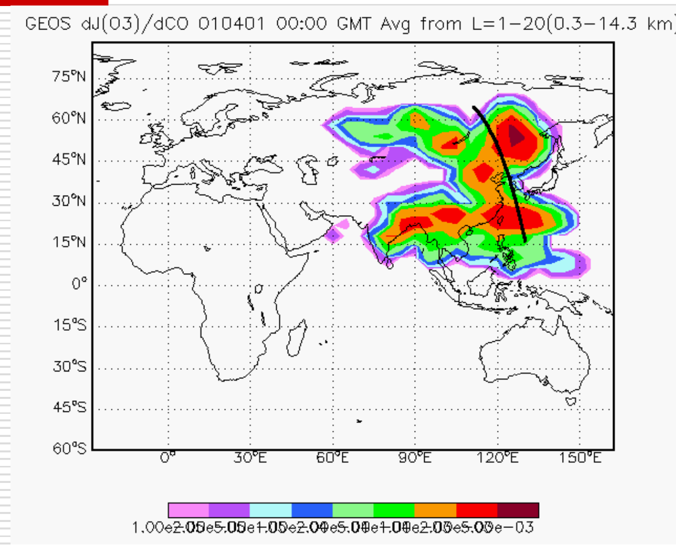

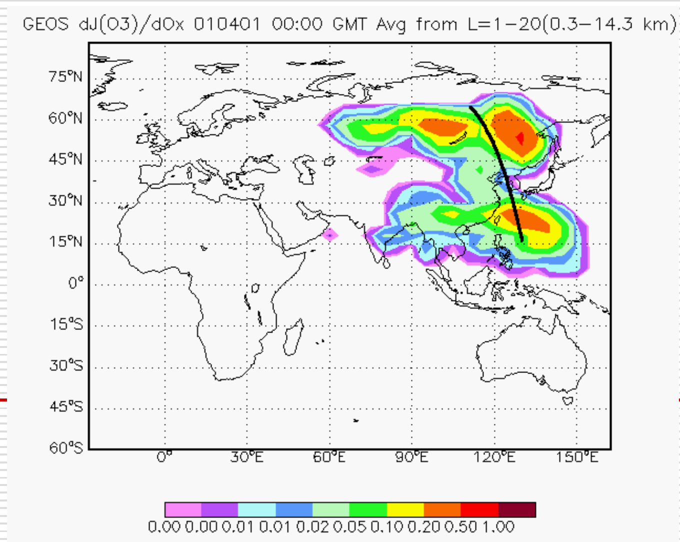

4-D Variational Data Assimilation

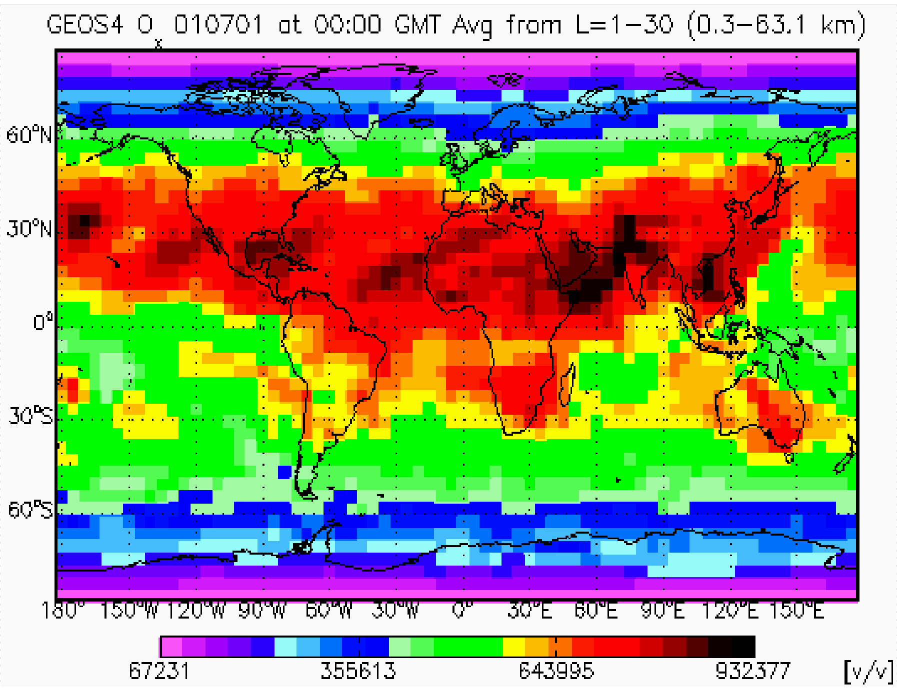

Difference between perturbed and reference concentration ($c_p^0 - c_0^0$)

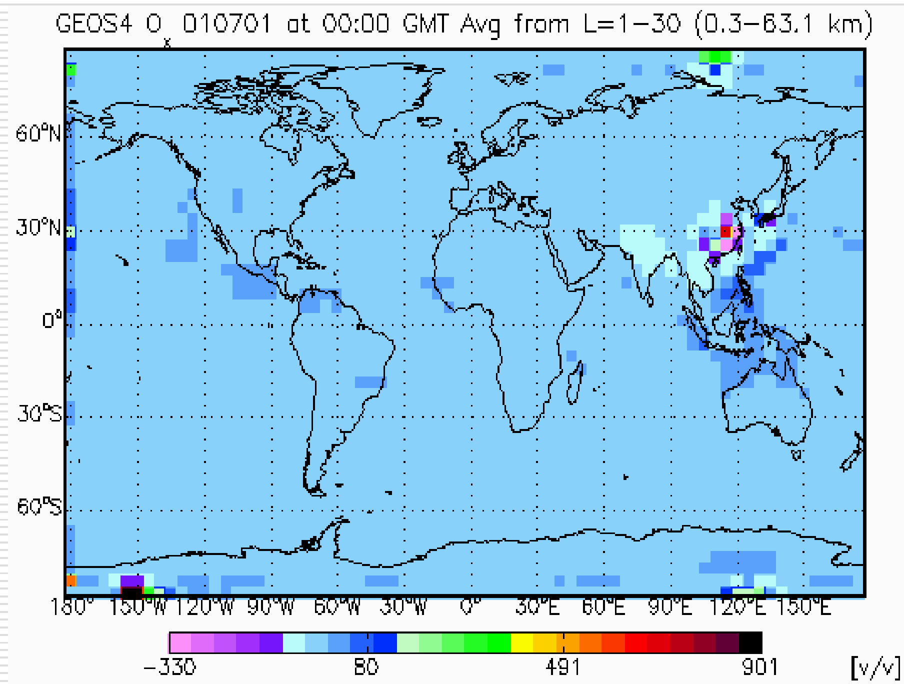

Difference between optimized and reference concentration ($c_p^0 - c_0^0$)

4-D Variational Data Assimilation

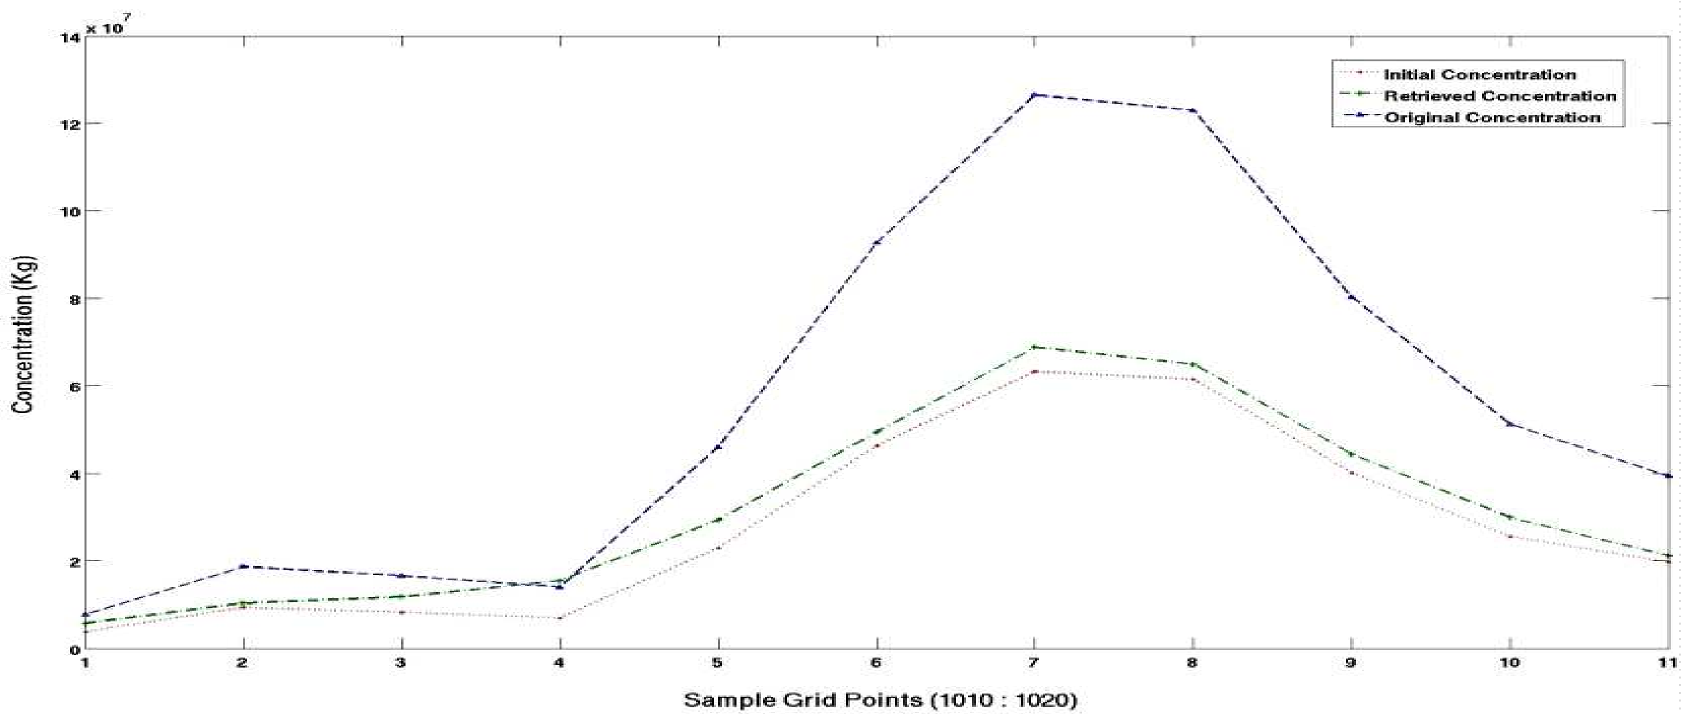

Plot of correction in the initial concentration for twin experiment run over 2006 summertime GEOS-Chem data for 3 days with TES profile retrievals generated synthetically

4-D Variational Data Assimilation

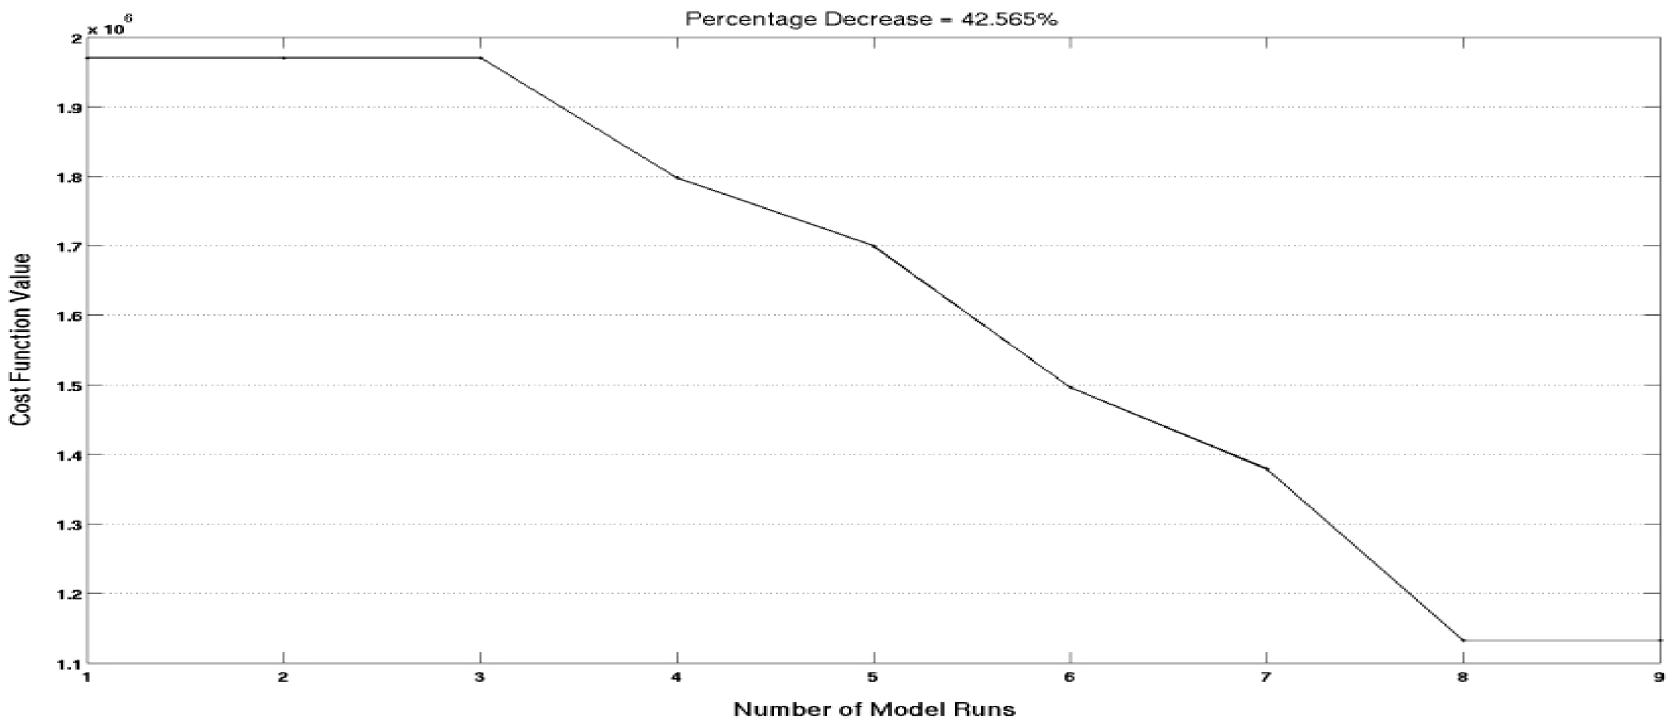

Plot of decrease in the cost function with respect to model runs for twin experiment run over 2006 summertime GEOS-Chem data for 3 days with TES profile retrievals generated synthetically

KPP Chemistry using CUDA

- Passed consistency test for 1 grid block and 1 thread (within 0.1%)
- Device emulation mode works perfectly for any combination of grid blocks and threads, and for any number of grid cells
- On GPU card, successful up to 1 block with 32 threads – 32 grid cells in total
- Issue with more threads per block or more blocks

Consistency Results (CUDA KPP)

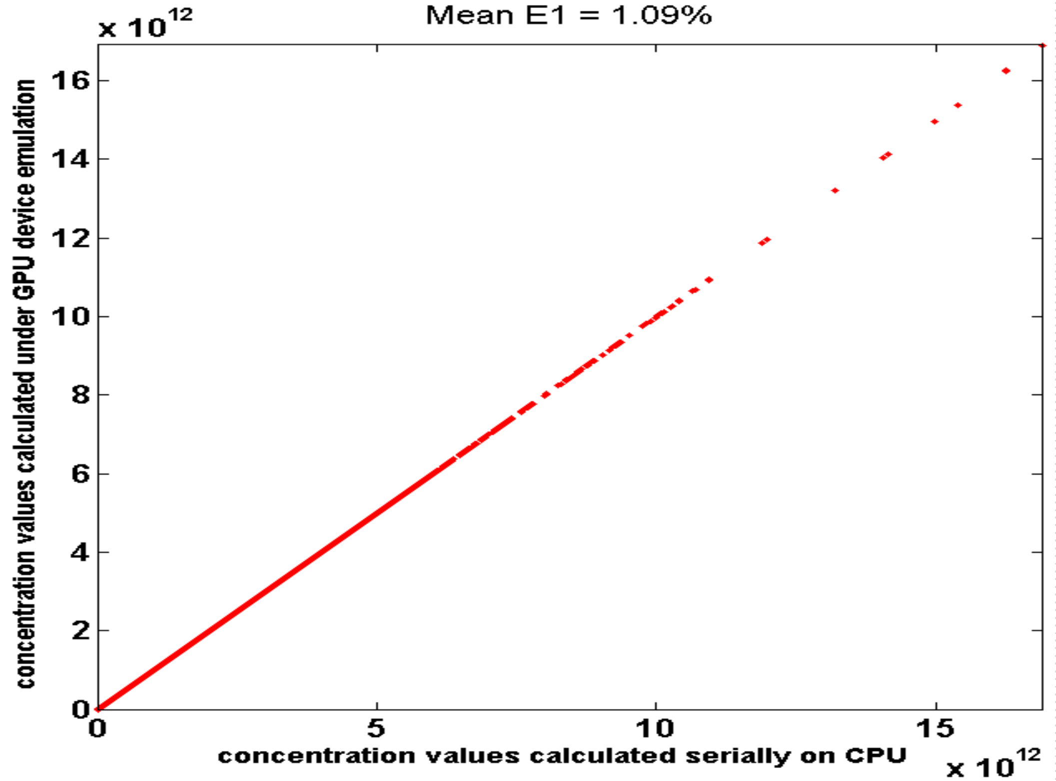

Scattered plot of concentration values calculated on CPU and GPU device emulation mode, 3200 grid cells

Consistency Results (CUDA KPP)

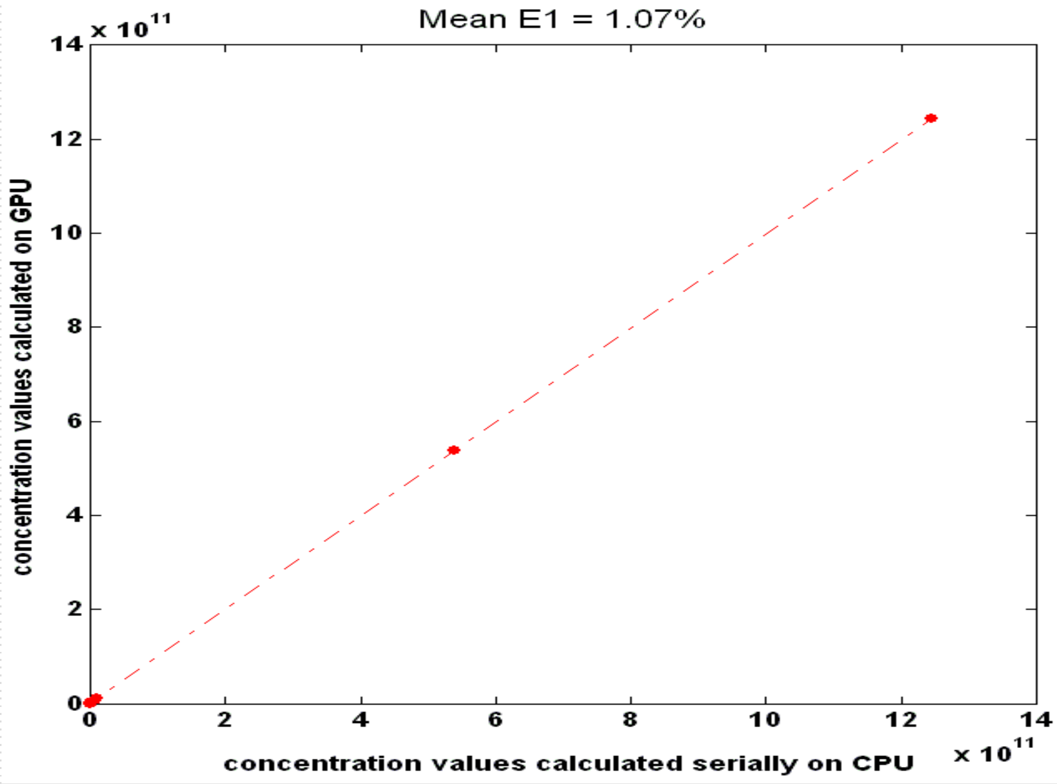

Scattered plot of concentration values calculated on CPU and GPU, 1 BLOCK and 32 threads, total grid cells 32

Conclusion

- Successfully constructed standard adjoint model of CMAQv4.5 – accepted by EPA
- Successfully developed an adjoint model of GEOS-Chem v7 – under final phases of delivery
- Parallelized GCv7 adjoint completely
- Added 4-D Variational data assimilation and sensitivity analysis capabilities to both the models
- Added Tropospheric Emission Spectrometer satellite observation operator and it's adjoint to include real data observations
- Created CUDA chemistry as part of the proposed activity

Future Work

- Data assimilation and sensitivity analysis using real data
- 4-D Variational approach against sub-optimal kalman filter
- Resolve issue with multiple blocks on GPUs
- Analyze performance of hybrid CPU and GPU approach against other high performance machines
- Submit the deliverables

Questions