Distributed Data Reduction and Analysis: An Overview with applications to Californian Climate and Energy Demand

Charlie Zender <zender@uci.edu>
Department of Earth System Science
University of California at Irvine

Presented to:
Calit2 SURF-IT Program
University of California at Irvine
August 23, 2005

Many Thanks to:
Mike Brown, Brian Bush, Chris Knox, Frank Wessel

Presentation PDF on the Web at:
http://dust.ess.uci.edu/smn/smn_ddra_surfit_2005.pdf

Abstract

Distributed Data Reduction and Analysis (DDRA) refers to the task of crunching datasets stored in separate physical or network spaces. DDRA has emerged as both an opportunity and chokepoint for collaborative research in high performance computing applications. For example, Climate simulations prepared for the Fourth Assessment Report (AR4) of the Intergovernmental Panel on Climate Change (IPCC) reside on a distributed network of storage archives known as the Earth System Grid (ESG). I will describe how we use DDRA techniques to get the data from the ESG and use it to characterize the envelop of future Californian climate contained within these datasets. Our goals are two-fold: 1. To quantify the Californian climate expected under a variety of IPCC forcing scenarios. 2. To benchmark, characterize, and reduce bottlenecks encountered in DDRA of geophysical datasets.

1. Distributed Data Reduction & Analysis (DDRA)

Challenge: How to optimize evaluation of Equation (1) when $\bar{x}_j$ depends on input elements x_i distributed across a network?

$$\bar{x}_j = \frac{\sum_{i=1}^{i=N} \mu_i m_i w_i x_i}{\sum_{i=1}^{i=N} \mu_i m_i w_i} \tag{1}$$

Strategy: Extending the NCO netCDF Operators (Unidata, 2004; Zender, 2004a), scientific analysis middleware, to accelerate DDRA (Zender, 2004b) by exploiting MPI-Grid and OPeNDAP (OPeNDAP, 2004), a discipline-neutral, network-transparent data accesss protocol running across the OptIPuter (Smarr et al., 2003), a high-bandwidth network.

Proof-of-Concept: The grid-enabled Scientific Data Operator (SDO) toolkit is applied to efficiently analyze and intercompare climate change simulations for the next Intergovernmental Panel on Climate Change (IPCC) assessment report (AR4). This tests DDRA on many TB of climate data stored across the Earth System Grid (LBNL, NCAR, and ORNL).

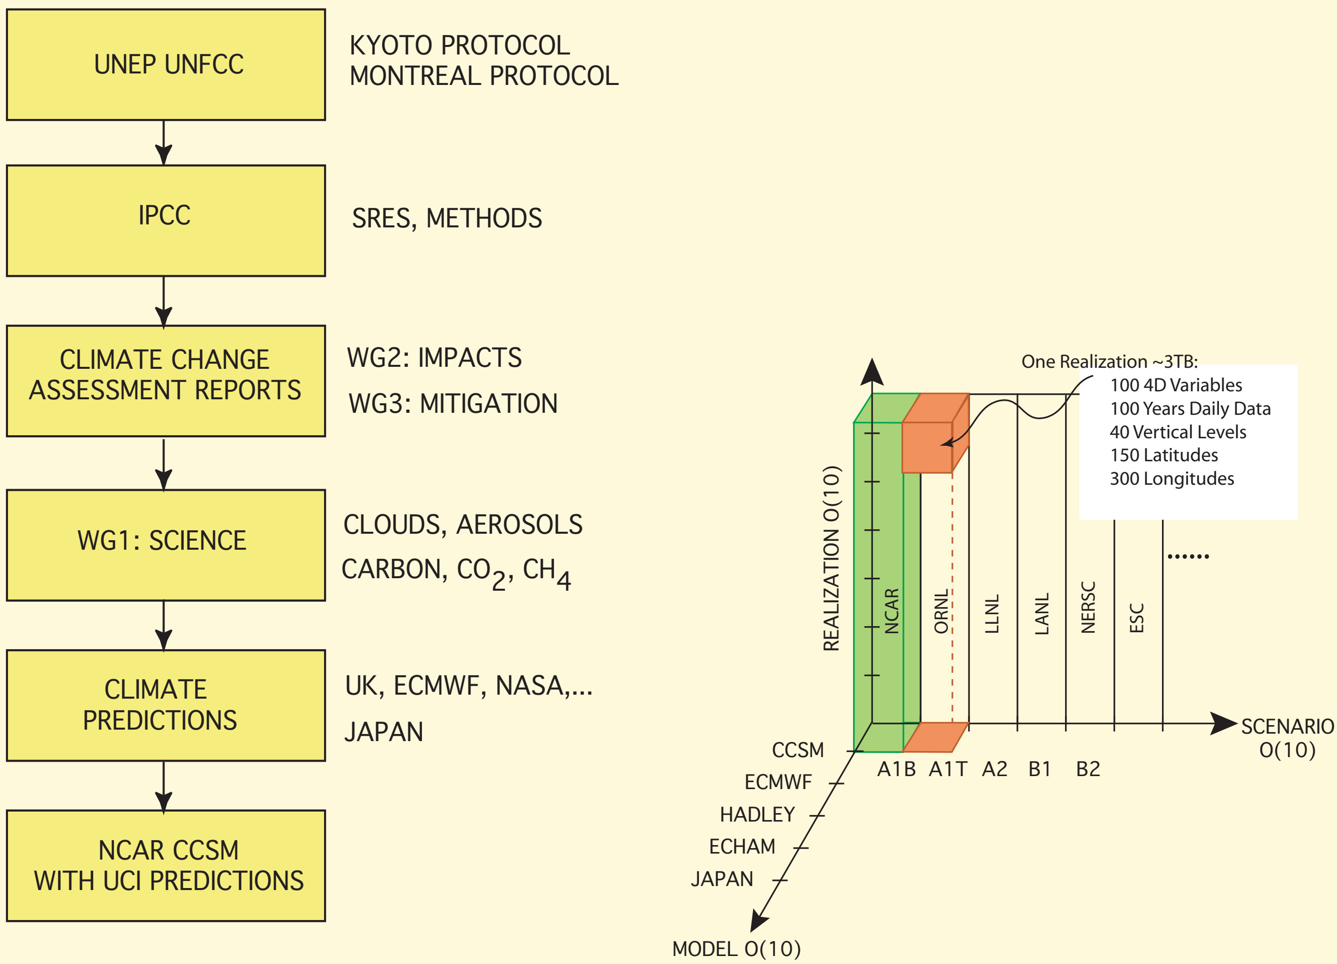

Figure 1: Distributed Data Reduction & Analysis of IPCC simulations (*Figure*: C. Zender & F. Wessel)

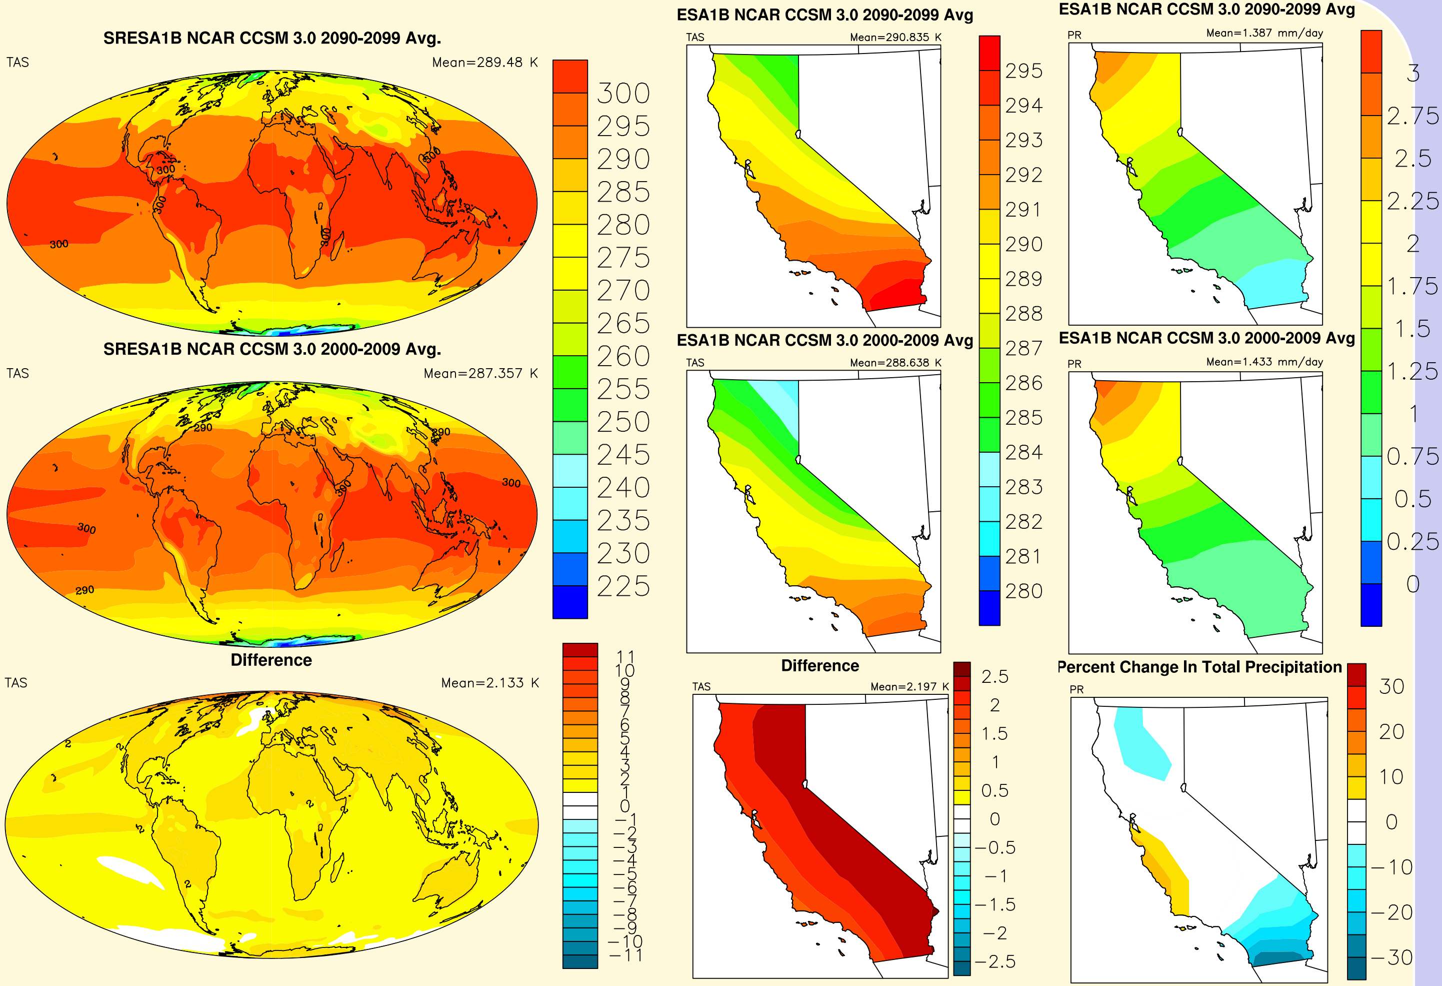

Figure 2: CCSM-predicted Global and California decadal mean temperature and precipitation from (a) 2090–2099, (b) 2000–2009, and (c) difference under SRESA1B 720 ppm CO_2 stabilization scenario.

2. DDRA Examples

Consider `operator`, which manipulates (e.g., adds, averages, concatenates) data:

1. Local Data Reduction and Analysis (LDRA):

```
operator file1 file2 ... fileK fileout
```

2. Distributed Data Reduction and Analysis (DDRA):

```
operator http://server1/file1 http://server2/file2 ... \
         http://serverK/fileK fileout
operator --openmp_thread_number=N ... [input_files] ... fileout
operator --mpi_processes_number=N ... [input_files] ... fileout
```

Performance of DDRA depends on many factors including: filesystem localization, network bandwidth and latency, arithmetic complexity and computability, server-side data reduction and analysis (SSDRA) …

3. Software Acceleration

We are benchmarking three types of DDRA acceleration:

1. Symmetric Multi-Processing (SMP) Parallelism (OpenMP library):
   Simultaneous execution on all available CPUs in single-node.
2. Single Program Multiple Data (SPMD) Parallelism (MPI library):
   Simultaneous execution on all available CPUs across multiple nodes.
3. Parallel I/O (MPI-IO library):
   Removes write bottleneck. Works in tandem with SMP or MPI.

4. NCO Code Examples

Symmetric Multi-Processor (SMP) parallelism with OpenMP:

```
#ifdef _OPENMP
#pragma omp parallel for default(none) private(idx) shared(dbg_lvl,fl_in ... )
#endif /* !_OPENMP */
for(idx=0;idx=var_nbr;idx++){ /* While un-processed variables remain... */
   nco_var_avg(var[idx]); /* ...process variables in parallel... */
```

Single Program Multiple Data (SPMD) parallelism with MPI:

```
while(1){ /* While work remains... */
  /* Send WORK_REQUEST */
  wrk_id_bfr[0]=proc_id;
  MPI_Send(wrk_id_bfr,wrk_id_bfr_lng,MPI_INT,mgr_id,WORK_REQUEST,MPI_COMM_WORLD);
  /* Receive WORK_ALLOC */
  MPI_Recv(info_bfr,info_bfr_lng,MPI_INT,0,WORK_ALLOC,MPI_COMM_WORLD,&mpi_stt);
```

Shell Scripts:

```
for mdl in 'cccma_cgcm3_1 cccma_cgcm3_1_t63 cnrm_cm3 csiro_mk3_0 \
gfdl_cm2_0 gfdl_cm2_1 giss_aom giss_model_e_h giss_model_e_r \
iap_fgoals1_0_g inmcm3_0 ipsl_cm4 miroc3_2_hires miroc3_2_medres \
miub_echo_g mpi_echam5 mri_cgcm2_3_2a ncar_ccsm3_0 ncar_pcm1 \
ukmo_hadcm3 ukmo_hadgem1'; do
        ncwa -R -O -D 3 -m msk_rgn -w area -a lat,lon \
             -p ftp://climate.llnl.gov/$scn/atm/mo/$var/$mdl/run$run \
             -l $DATA/$scn/atm/mo/$var/$mdl/run$run $var_A1.nc \
             $scn_$mdl_$run_$var_$yyyymm_$yyyymm.nc
```

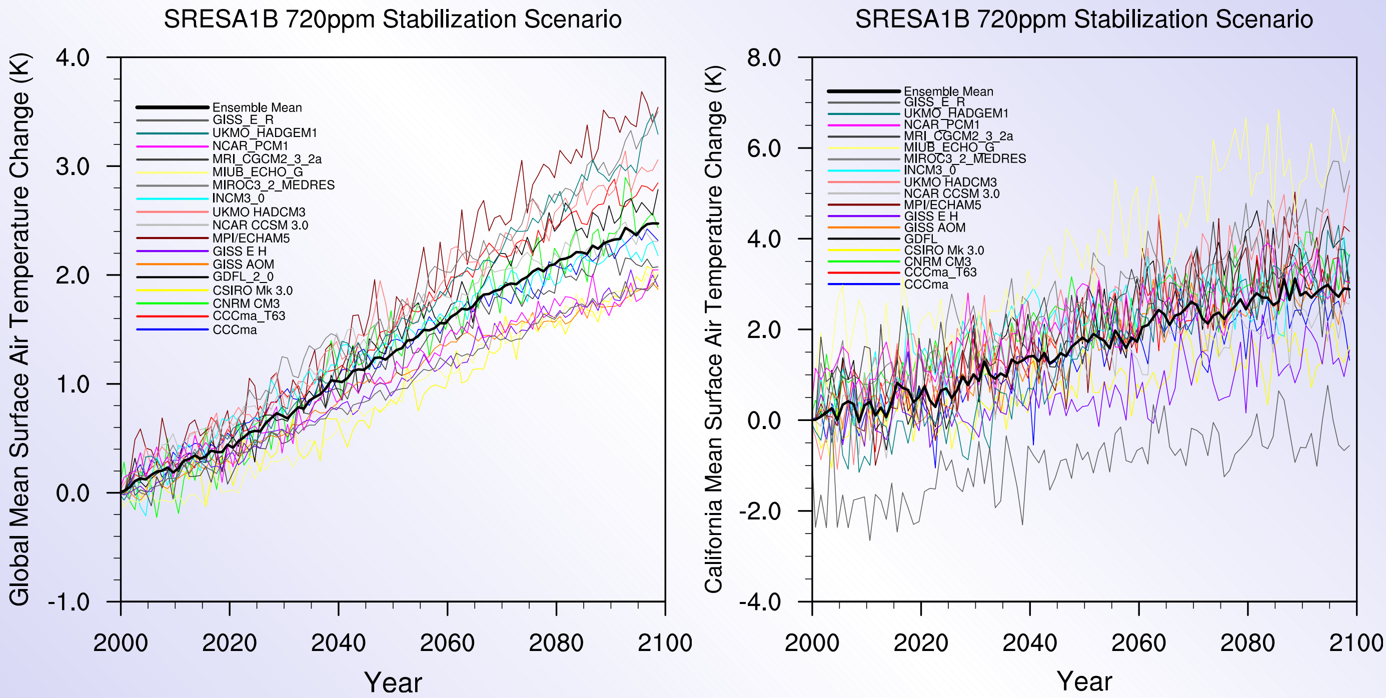

Figure 3: Predicted Global (left) and California (right) annual-mean temperature from 2000–2099 under SRESA1B 720 ppm CO_2 stabilization scenario. Temperature scales differ.

Figure 4: HIPerWall displaying ensemble of IPCC AR4 Model results (Photo: Chris Knox)

5. Energy Demand

1. The largest economic impact of climate change on California is expected be increased power demand to cool buildings and homes (Electric Power Research Institute, 2003).
2. Energy demand increases about 2.0% $^{\circ}\mathrm{F}^{-1}$, with non-linear effects such as infrastructure upgrades
3. Climate trends are probably the fourth most important determinant of power demand trends in California, following population, industrialization, and technology. However, the $\sim 5.0\,^{\circ}$F warming predicted by 2100 is not accounted for in California's Integrated Energy Planning Reports.
4. How to plan for climate-related energy demand trends?

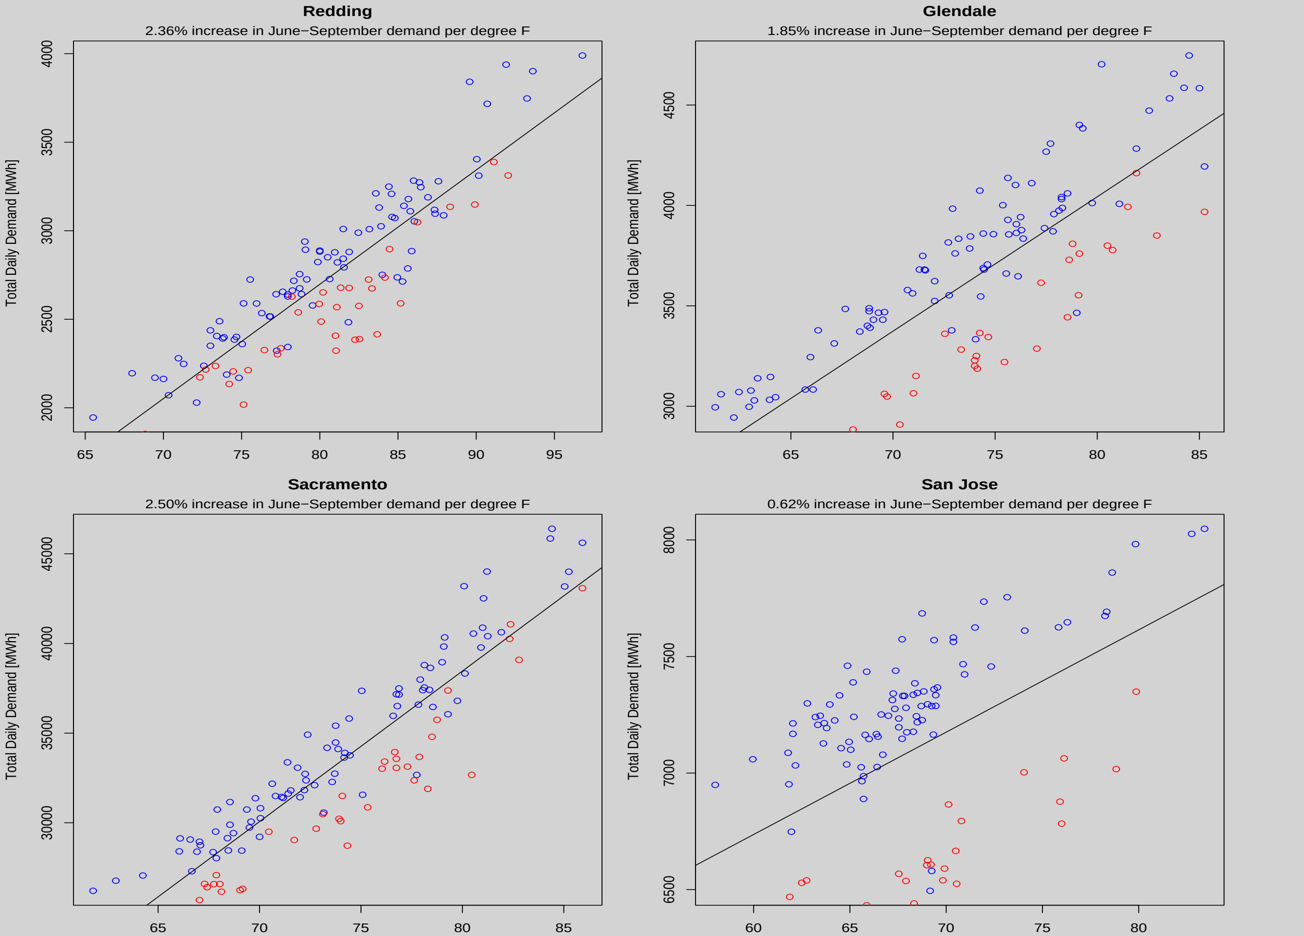

Figure 5: Correlations of 2003 Daily energy demand and temperature in (a) Redding, (b) Glendale, (c) Sacramento, and (d) San Jose. Source: Brian Bush, LANL

6. Future Work

1. Use DDRA on IPCC simulations to characterize expected Californian climate change (M. Brown)
2. Intermediate Terascale Visualization and Analysis (DVA), e.g., HIPer-Wall (C. Knox, F. Kuester, S. Jenks)
3. Move DDRA from command-line to web services? (S. Jenks)
4. Convolve Californian climate change trends with empirically derived climate sensitivity of energy demands to estimate climate-driven energy demand trends in every CA county (B. Bush)
5. Server-side Data Reduction and Analysis (MPI-Grid, OPeNDAP): Distributes intermediate computations to data locale (H. Mangalam)

7. Where to Learn More

References

Electric Power Research Institute, 2003: Global climate change and California: Potential implications for ecosystems, health, and the economy. Consultant Report 500-03-058CF, California Energy Commission, Sacramento, CA.

OPeNDAP, 2004: Open-source Project for a Network Data Access Protocol. http://www.opendap.org.

Smarr, L. L., A. A. Chien, T. DeFanti, J. Leigh and P. M. Papadopoulos, 2003: The OptIPuter. *Comm. ACM*, **46**(11), 58–67.

Unidata, 2004: *Network Common Data Format*. Boulder, CO, http://www.unidata.ucar.edu/packages/netcdf.

Zender, C. S., 2004a: NCO User's Guide. "http://nco.sf.net/nco.pdf".

Zender, C. S., 2004b: Scientific data operators optimized for distributed interactive and batch analysis of tera-scale geophysical data. “http://dust.ess.uci.edu/prp/prp_sei/prp_sei.pdf”.