The PSAS Software Workshop, October 26 1998
Short glossary of terms

Physical-space Statistical Analysis System (PSAS)
w^f: forecast vector of n gridpoint values
w^a: analysis vector of n gridpoint values
Analysis grid: currently the analysis grid has resolution $2.5° \times 2° \times 18$ pressure levels, therefore $n \approx 10^6$
w^o: observation vector of p observation values
For a typical synoptic (6 hours) period $p \approx 10^5$
Innovation a.k.a. observed minus forecast (OMF): $w^o - Hw^f$
Innovation matrix: $M = HP^fH^T + R$; sometimes Jing and Jay use the generic C to represent one of the matrix inner factors (see factored operator form below)
Innovation equation: $(HP^fH^T + R)x = w^o - Hw^f$
Analysis equation: $w^a - w^f = P^fH^Tx$
Factored operator form: the breakdown of P^f and R into the sum of factored operators (see today's handouts, viz, Guo et al. 1998, and Larson et al. 1998 Eqs. 11 and 12)
Regions a.k.a. kr regions: icosahedral spherical segments triangulated to make 80 regions that cover the sphere
kr/kt segment: a subvector whose elements have the same values of region index (kr) and data type (kt)
kr/kt blocks: sub-blocks of the matrices $(HP^fH^T + R)$ or P^fH^T (or their corresponding matrix factors) corresponding to element-wise correlations between the members of kr/kt segments or members of kr/kt segments and analysis grid points
Blocking: aggregation of matrix blocks in an algorithm
Rank, pid, node number: the processor ID number in a distributed parallel algorithm
pe: processing elements (pes)
Replication: reproducing a data structure (e.g., vector segment) on all processors in a distributed-memory algorithm
Partial replication: reproducing a data structure (e.g., vector segment) on some but not necessarily all processors in a distributed-memory algorithm
MPI_Allgatherv: gathers data from all tasks and deliver it to all
MPI_Reduce_scatter: combines (sums) values and scatters the results
GEOS DAS: Goddard Earth Observing System Data Assimilation System
GEOS-1: version of GEOS DAS that used Optimal Interpolation (OI) analysis
GEOS-2: version of GEOS DAS that used the PSAS analysis
Versions of the PSAS: Early PSAS (1991-1995), GEOS-2 PSAS (1995-), JPL (prototype) PSAS
GEOS-3: the distributed-memory (MPI/shmem) version of GEOS (1999)

Short summary of numbers

Number of analysis gridpoints $n \approx 10^6$
Number of observations for a typical synoptic (6 hours) period $p \approx 10^5$
On the Origin 2000 machine Turing, the interprocessor bandwidth for point-to-point MPI communications is about 140 Mbytes/sec, while the latency is about 10 microsec.
Lookup tables in GEOS-2 PSAS use currently 14 MWord of memory
One analysis (PSAS) on 8 processors of the Origin 2000 (195 MHz) takes 25 minutes