

AD-A251 242



(2)

PL-TR-92-2011

An Enhanced Global Spectral Model

Thomas Nehrkorn
Ross Hoffman
Jean-François Louis
Marina Zivkovic



Atmospheric and Environment Research, Inc.
840 Memorial Drive
Cambridge, MA 02139-3794

27 February 1992

Final Report
September 1989 - November 1991

Approved for public release; distribution unlimited




PHILLIPS LABORATORY
AIR FORCE SYSTEMS COMMAND
HANSCOM AFB, MASSACHUSETTS 01731-5000


92-13366




92 5 19 035 +

This technical report has been reviewed and is approved for publication.


DONALD C. NORQUIST
Contract Manager


DONALD A. CHISHOLM, Chief
Atmospheric Prediction Branch


ROBERT A. MCCLATCHEY, Director
Atmospheric Sciences Division

This document has been reviewed by the ESD Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS).

Qualified requestors may obtain additional copies from the Defense Technical Information Center. All others should apply to the National Technical Information Service.

If your address has changed, or if you wish to be removed from the mailing list, or if the addressee is no longer employed by your organization, please notify PL/TSI, Hanscom AFB, MA 01731-5000. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 27 February 1992		3. REPORT TYPE AND DATES COVERED Final September 1989 - November 1991
4. TITLE AND SUBTITLE An Enhanced Global Spectral Model			5. FUNDING NUMBERS PE 63707F PR 2688 TA 04 WU JB	
6. AUTHOR(S) T. Nehr Korn, R. N. Hoffman, J-F. Louis, M. Zivkovic			Contract F19628-89-C-0112	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Atmospheric & Environmental Research, Inc. 840 Memorial Drive Cambridge, MA 02139-3794			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Phillips Laboratory Hanscom AFB, MA 01731-5000 Contract Manager: Donald Norquist/GPAP			10. SPONSORING/MONITORING AGENCY REPORT NUMBER PL-TR-92-2011	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report describes the development of a vectorized, multiprocessing global spectral model (GSM) with enhanced physical parameterizations. The starting point was the Phillips Laboratory GSM, in its GL90 version, containing an enhanced suite of physical parameterizations. The latitude tasking scheme for multiprocessing the loop over latitude in the calculation of the spectral tendencies and adjusted model variables was implemented, using the general truncation version of the hydrodynamics code. Wavenumber calculations were vectorized over wavenumber and multiprocessed over vertical level. All gridpoint calculations were vectorized over longitude, and the physics packages were brought into closer compliance with plug-compatibility rules. Speedups due to the optimization were demonstrated in single- and multiproc-essing timing tests on a dedicated Cray 2 and Cray Y-MP. The enhanced physics GSM was evaluated in forecast tests, and contrasted with the simpler physics GSM used at GWC. Minimization of errors in the computation of the horizontal pressure gradient force was investigated, using a perturbation temperature instead of full temperature in the integration of the hydrostatic equation. Several schemes were tested in an idealized model atmosphere.				
14. SUBJECT TERMS Numerical weather prediction, multiprocessing, vectorization, global spectral model			15. NUMBER OF PAGES 84	
			15. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

TABLE OF CONTENTS

	Page
1. Introduction.....	1
2. Description of model code.....	3
2.1 Hydrodynamics and overall design.....	5
2.2 Radiation	6
2.3 Planetary Boundary Layer.....	6
2.4 Adjustment Physics	7
3. Timing Results.....	8
3.1 Single-Processor Timings	8
3.2 Multi-Processor Timings	12
3.3 Overall Speedups.....	14
4. Forecast Results	15
5. Minimization of pressure gradient force errors.....	27
5.1 Pressure-gradient force errors.....	27
5.2 Error minimization.....	28
5.2.1 Schemes	31
5.2.1.1 The PG - scheme	31
5.2.1.2 The MPG -scheme.....	32
5.2.1.3 The null scheme - an isothermal case....	32
5.2.2 Code modification.....	32
5.2.2.1 Model equations	34

5.2.2.2 Code changes.....	35
5.3 Initial Data.....	36
5.4 Tests.....	37
5.4.1 Scheme comparison.....	38
5.4.2 Vertical distribution of the model levels.....	45
5.4.3 The model diffusion.....	46
5.5 Conclusions.....	48
6. Summary.....	49
7. References.....	50
Appendix A: User's Guide.....	A-i
Appendix B: List of Variables.....	B-i

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist.	Avail and/or Special
A	



1. Introduction

This report describes the work performed by AER under contract F19628-89-C-0112 for the US Air Force Phillips Laboratory, Geophysics Directorate (hereafter referred to simply as PL), entitled "An Enhanced Global Spectral Model", for the entire period of the contract, 27 September 1989 to 27 November 1991. The purpose of the work was the development of a vectorized, multiprocessing global spectral model (GSM) with enhanced physical parameterizations. The starting point was the Phillips Laboratory (then called Geophysics Laboratory, or GL) GSM, in its GL90 version.

The global spectral model developed by PL has evolved from the GSM used at NMC as described in (Sela, 1980). The hydrodynamics were completely redesigned (Brenner et al., 1982, 1984). This model has been used by the Air Force Global Weather Central (GWC) with the physics routines taken almost intact from NMC (circa 1983). More recently, an enhanced suite of physical parameterizations has been developed by different investigators and integrated by PL personnel. The new physics routines consist of a coupled boundary layer and soil model (Mahrt et al., 1984, 1987), a radiation parameterization (Liou et al., 1984; Ou and Liou, 1988), and a modified Kuo convection parameterization (Soong et al., 1985 and Norquist and Yang, 1990). The new physics routines have been used in a research mode in a number of studies by PL personnel which demonstrated their potential for improving forecast skill.

As originally implemented, the new physics package required significantly more CPU time than the baseline GWC physics. Thus, before the potentially better forecast skill could be realized in the operational environment, the execution time of the enhanced global spectral model had to be reduced.

Implementation of the code on the Cray 2 offers two major avenues for reducing wall-clock execution time: vectorization and parallelism. Vectorization reduces the CPU time (and, hence, the wall-clock time) spent by a single processor by performing identical (or similar) operations on a number of storage locations in a pipeline fashion. This technique was

pioneered by Cray Research, and extensive support tools for vectorizing scalar codes are now available as part of compilers on Cray and other vector machines. To take full advantage of the potential speed-ups, however, it is necessary to take into account vectorization considerations during the design phase of program development.

A more recent trend in computer technology is the construction of machines composed of a number of processors linked together, allowing multiple simultaneous computations. Such multiprocessing computers contain two or more, in some cases tens of thousands, of individual processing units. These processors may operate in lock step or they may be nearly autonomous. The former case, single instruction multiple data (SIMD), is particularly relevant for image processing applications, while the latter case, multiple instruction multiple data (MIMD), is of more general interest and includes the Cray 2 and Y-MP. Multiprocessor computers give a substantial increase in computing speed for large complex numerical problems. Thus, they provide an opportunity to reduce significantly the computing time for GCM studies and NWP. However, a special effort is required to exploit this opportunity. According to the recommendations of GARP Special Report No. 43 (WMO, 1985), "[I]ncreasing computer power and changes in computer architecture...[have] always played a major part in determining integration techniques and will continue to do so if advances in computer technology are to be fully exploited. Therefore, study to determine the most efficient or appropriate algorithms will continue to be necessary, taking into account, for example, the future use of multiprocessors with a very large storage. Development costs of such optimized computer codes will be high, and they should be designed from the outset to be widely usable..."

As an add-on to the original contract, we investigated ways of minimizing errors in the computation of the pressure-gradient force in the sigma-coordinate system. This work concentrated on the use of perturbation temperature in the integration of the hydrostatic equation, and was prompted by work of Simmons and Chen (1991).

The design of the optimized version of the GSM is described in detail in the interim technical report of this contract (Nehrkorn et al., 1990;

referred to as I in the following). The main points of the design, and differences in the final implementation from the original design, are discussed in Section 2. Detailed user documentation in the form of a user's guide and an annotated list of model variables are provided as Appendix A and B, respectively; the appendices are also included as plain text files in machine-readable form, with the delivered model code. Section 3 describes the results of timing tests of the original and optimized version of the GSM with enhanced physics, along with a version of the GSM with simple physics (similar to what is used operationally at the US Air Force Global Weather Central - GWC). Section 4 contains the results from 4-day forecasts, using the GSM with enhanced and simple physics. Section 5 describes results of the study to minimize errors in the computation of the horizontal pressure gradient force. A summary and concluding remarks form Section 6, and Section 7 contains the references.

2. Description of model code

The starting point for our optimization effort was the GL90 version of the GSM. The physics modules of this version consisted of the radiation module with a one-cloud deck structure, the PBL scheme with globally uniform soil and vegetation characteristics, and the Kuo convection scheme. Simultaneous with our recoding for computational efficiency, testing and tuning of the physics packages were carried out by PL personnel, and we were able to incorporate some of the resulting changes into our optimized design. In particular, we used a three-deck version of the radiation for optimization, and this was incorporated both into the "original" and optimized code used in the forecast and timing tests. Additional minor changes in the Kuo convection and PBL scheme were also incorporated at the same time. We refer to this version of the optimized model code as version 5.1 (all modules of the optimized code are archived with the UNIX Source Code Control System (SCCS), and all are identified by the SCCS sequence number 5.1).

Further modifications were made to the radiation parameterization that could not be incorporated into the model in time for the forecast and timing tests. These changes involved the use of latitudinally varying climatologies of ozone, and replacement of the Geleyn cloud parameteriza-

tion by the scheme of Slingo (1987). The latter also required some minor code changes to the convection and hydrodynamics modules. For our work with the GSM under another contract, these changes were incorporated into the optimized GSM, and the resulting model version was archived as version 5.2. Both versions of the code have been delivered, and the user's guide document (Appendix A) is applicable to both versions.

There have been other improvements made to the physics modules by PL personnel that have not yet been incorporated into the optimized model code. Most notably, these are the addition of a gravity wave drag scheme, the replacement of the Kuo convection scheme by a mass flux scheme, and the use of geographical databases of soil and vegetation characteristics.

Our changes to the original code were guided by two general strategies.

Enhance clarity. Clear, direct coding is best. Coding which makes special use of a particular machine architecture and compiler is in the long run counterproductive. Eventually, and probably sooner than later, the compiler will be updated or the code will be transported to other machines. Generally, as time goes on, compilers get more and more capable and what seemed to be necessary last year is now a hindrance. Complicated coding strategies to squeeze the maximum horsepower out of the CRAY2 will be restricted to very specific and isolated procedures. For the most part these procedures will be standard library routines for which vanilla versions are readily available. An example is the FFT software.

Avoid impacts. Do not make changes which will have a big impact on the results. Changes which result in small changes in the tendencies, which are not less correct are fine. Some differences in the calculated tendencies are unavoidable. It is expected that different versions of the model, or the same version compiled differently will usually result in a different order of arithmetic and hence in different round off errors. On the other hand changes which result in large differences in the tendencies, even though they may be as correct as the original, should be avoided. For example, one could argue that a speedup is possible by using a T150 truncation instead of a R120 truncation. However, demonstrating that this is

indeed the case, would require a substantial suite of test cases, well beyond the scope of the present effort.

2.1 Hydrodynamics and overall design

The latitude tasking scheme described in I was implemented, using the general truncation version of the hydrodynamics code developed under a previous contract. The multitasking was implemented through Cray microtasking directives (these are treated as comments if Cray multitasking is deactivated, or if other compilers are used). Separation of data into shared (global) and private (i.e., local to each latitude task) was accomplished with Cray Fortran TASK COMMON statements - on single-processor computers, these would be replaced by simple COMMON statements, and on other multi-processing computers with language extensions analogous to the Cray TASK COMMON statement.

The general truncation version (documented in Kaplan et al., 1985) supports any pentagonal truncation, including the frequently used rhomboidal and triangular truncations. In addition, computations may be performed over only one hemisphere or a section of the globe. Array dimensions are given in terms of PARAMETER statements. The parameter statements and the common blocks are stored in include files, which are incorporated into the code at compile time. Edit symbols are retained, however, to allow to choose between the hemispheric and global versions of the code, and to use double precision for certain, sensitive calculations (useful for use of the code on 32-bit machines; see Nehrkorn, 1990). This version of the code also makes use of Hoffman's (1985) semi-implicit time stepping scheme, which uses full temperature (as opposed to perturbation temperature) throughout. This does not preclude introducing a basic state and perturbation temperature for the integration of the hydrostatic equation. In general, the basic state temperature used for the hydrostatic equation and that used for the semi-implicit time scheme will be different. Therefore it is best to use full temperature throughout and introduce linearizations where needed. In our study described in Section 5, the basic state temperature is a function of pressure (which depends on latitude, longitude, and time for each sigma layer), and the hydrostatic calculation is integrated in grid point space; the semi-implicit scheme still uses a

linearization about a globally uniform basic state temperature as a function of sigma.

The input/output was adapted to the PL practice, with minor changes and corrections (see Appendix A for details).

The gridpoint calculation in the hydrodynamic part of the code (in subroutine NLPROD) were completely recoded such that longitude was the innermost loop index, resulting in much improved vectorization.

2.2 Radiation

The GL90 version of the radiation, modified for a three-deck cloud structure, was the starting point for this work. The code was redesigned for plug-compatibility and vectorization as described in I for version 5.1 of the GSM. In particular, algorithms were redesigned to allow efficient vectorization, with longitude as the inner-most loop index. The holding file of radiative heating rates (needed for time steps when radiation computations are not performed) has been replaced by in-core storage in a global array. On the Cray 2, this did not lead to memory problems, even for the high-resolution runs (rhomboidal truncation at wavenumber 120, R120). However, we have included code to store these data in a direct access file which could be accessed simultaneously by the different latitude tasks.

As discussed above, a version of the radiation code was delivered (as version 5.2 of the model code) that included several enhancements: global climatology of ozone, and Slingo (1987) cloud parameterizations. This version of the code was not used in the forecast and timing tests reported below, however.

2.3 Planetary Boundary Layer

We implemented the design proposed in I by reordering all loops such that longitude was the innermost index. Vertical loop limits that were dependent on longitude had to be replaced by flags that indicated regions where calculations would or would not be performed. Similarly, surface computations that depend on the surface type (snow, land, or ocean) are now handled by flags. Computations that were unnecessary

were eliminated (commented out) - most of these involve calculations of boundary layer clouds, which were not used anywhere in the original code.

Diagnostic calculations were isolated in a separate entry point, which is called after all the latitude task computations are completed. Values needed for diagnostic calculations are stored in internal common blocks. In the original code, the diagnostic computations used algorithms that are not identical to those used in the actual tendency calculations; we have noted these inconsistencies, but have not made any substantive changes to the diagnostic computations.

The plug-compatibility rules are adhered to for the most part, except that we still use common blocks (in `ground.blk`) to pass values of boundary data sets between the PBL and the rest of the GSM.

2.4 Adjustment Physics

The adjustment physics consists of three separate, plug-compatible modules: dry adiabatic adjustment, large-scale precipitation, and Kuo convection. The cover routine for these routines performs spectral transforms to and from gridpoint space, and was adapted from the general truncation version of the GSM.

The large-scale precipitation module was already vectorized in the original code, and we only changed the interface to the cover routine in order to bring it into compliance with the plug-compatibility rules.

The dry adiabatic adjustment was redesigned for vectorization over longitude, as described in detail in I (the final version of the code differs slightly from the listing reproduced in I, however). The interface was also changed for plug-compatibility.

We had originally developed a vectorized design for the Kuo convection (as described in I), but did not pursue its implementation and testing to completion when it became evident that the final physics package assembled by PL would not contain a Kuo scheme. Instead, we implemented the original code, with changes only to the interface and the initialization of physical constants.

3. Timing Results

3.1 Single-Processor Timings

Three versions of the GSM were used for timing tests. The first, denoted GWC in the following, consists of the optimized GSM code for the hydrodynamics and physics modules that emulate current GWC operational practice. The GWC physics does not contain any radiative parameterization, and a very simple drag-law type boundary layer parameterization. The adjustment physics consists of the same large-scale precipitation and dry adiabatic package as the full physics GSM, but the Kuo convection is disabled at most points through the use of high threshold values and various criteria that must be met before moist convection is allowed to take place. In our tests, we emulated this by using a dummy routine for Kuo convection. The second version is the original code of the full physics GSM, before any optimization. Finally, the third version is the optimized version of the full physics GSM. This version is identified by its SCCS sequence number as version 5.1. Thus, comparisons between the GWC and optimized timings will give an indication of the computational cost of the enhanced physics packages, while comparisons between the original and optimized code show the speedup due to optimization.

All timing tests were performed as history restarts from a previous 24-hour forecast of the GSM, to eliminate any spin-up that would distort the timing information. Each timing test consists of a three-hour forecast; since radiation is computed every three hours, timings can thus be extrapolated to multiples of this forecast length. All timing tests use the 18-layer structure used extensively by PL (e.g., Norquist and Yang, 1990). Table 1 contains the timings from our single-processor timing tests for the three model versions on a Cray 2, for a model resolution of rhomboidal 40 (R40) with 18 layers in the vertical. The differences in CPU time between these runs are summarized in table 2. All runs (except run 1c) were conducted in a dedicated mode, and CPU and elapsed (wall-clock) times are essentially identical.

Table 1: Cray 2 timings of single processor runs

Run number	Model version	Compiler optimization	CPU time	elapsed time
1a	original	vectorized	253.5	257.2
1b	optimized	vectorized	133.9	134.7
1c	GWC	vectorized	91.2	
2a	original	scalar	676.2	678.4
2b	optimized	scalar	525.4	527.3

Table 2: Cray 2 single processor run speedups

Run numbers	Description	speedup
1b - 1a	Effect of recoding for optimization	1.9
1c - 1b	Cost of advanced physics	1.5
1a - 2a	Vectorization of original code	2.7
1b - 2b	Vectorization of optimized code	3.9

It is evident from table 2 that there is about a two-fold speedup due to the optimization of the original code. Comparison of the timings for runs

with vectorization turned on and off indicates that this is due mostly to the better vectorization speedup in the optimized code (3.9 vs. 2.7), although there is also a slight speedup in the scalar run. The computational cost of the enhanced physics is only a 50% increase in CPU time for the optimized code (recall that the GWC code in run 1c uses the optimized hydrodynamics code). In current operational practice at GWC, only 12 layers are used in the vertical; if this is taken into account, the computational cost is approximately 120%.

A further breakdown of timings for the original and optimized model code is shown in table 3. The timings shown are derived from a repeat of runs 1a and 1b, in which the flowtrace option was enabled. The flowtrace utility on the Cray 2 results in estimates of CPU time spent in each subroutine (along with the number of times it is called). The numbers in table 3 were obtained by adding the CPU times of the appropriate subroutines (in cases where the same subroutine was called from different parts of the calculation, the total times were divided among the tasks). A comparison of the original and optimized timings shows that the major speedup stems from the optimization of the physics packages (radiation and pbl), with further speedups due to the adiabatic nonlinear products computation (nlprod), and more efficient use of the FFTs. As was discussed earlier, the Kuo scheme, which is the most important part of the adjustment physics, was not optimized, and there is consequently little change in the CPU time between the original and optimized code. The Legendre transforms indicate a speedup in the case of the Gaussian quadrature, and a slowdown for the Legendre sums, resulting in a slight overall speedup. The time stepping computation was also greatly accelerated, but it was only a minor contributor even in the original code.

Considering just the results of the optimized code from table 3, one can identify further areas of potential improvement. Aside from the existing Kuo scheme, which will be replaced by a different algorithm, the spectral transforms should be examined for further speedups, since they now consume more than half of the CPU time. The all-fortran version of the FFTs should be replaced by whatever optimized routines exist for a particular platform, and ways of rearranging the Legendre transforms to

Table 3. Flowtrace CPU times from 3-hour single-processor forecast on the Cray 2, for the original and optimized model code.

	original			optimized		
	sec	(% of total)	sec	(% of total)	sec	(% of total)
Calculation:						
Total:	250.000	(100.0 %)			132.000	(100.0 %)
1. Transforms:	101.313	(40.5 %)			79.719	(60.4 %)
1.1. FFTs:			39.542	(15.8 %)		21.107 (16.0 %)
1.2. Gaussian quadrature:			32.584	(13.0 %)		24.557 (18.6 %)
1.3. Legendre sums:			27.909	(11.1 %)		33.615 (25.5 %)
1.4. Pmns:			1.278	(0.5 %)		0.440 (0.3%)
2. Gridpoint calculations:	118.601	(47.4 %)			48.572	(36.8 %)
2.1. Radiation:			20.611	(8.2 %)		4.669 (3.5 %)
2.2. PBL:			43.947	(17.6 %)		4.249 (3.2 %)
2.3. Adjustment physics:			29.503	(11.8 %)		27.678 (21.0 %)
2.4. adiabatic nlprod:			24.540	(9.8 %)		11.976 (9.1 %)
3. time stepping:	29.946	(12.0 %)			3.192	(2.4 %)
4. Miscellaneous:	0.140	(0.0 %)			0.517	(0.4 %)

take advantage of optimized linear algebra routines should be considered. Other minor speedups may be possible throughout the code, and some have been identified in comment cards in the optimized physics packages. The results from these further modifications are bound to be machine-specific, however, and are only cost effective once a particular platform has been selected for extended use of the code.

3.2 Multi-Processor Timings

We conducted a number of other timing tests, all designed to test how well the optimized code can take advantage of multiple processors, and how its performance scales with increased resolution. The timings are shown in table 4; again, all runs were conducted on a dedicated Cray 2, except the 8-processor run, which was done on a dedicated Cray Y-MP. The R40 runs were conducted with the code compiled for multiprocessing, with 1, 2, and 4 processors. Comparison of the 1-processor run with the corresponding run in table 1 indicates that the overhead for multiprocessing is negligible (less than a 1% increase in elapsed time). The 2-processor run is about 1.87 times faster than the single-processor run, which corresponds to an efficiency of 93 %. Here efficiency is defined as the ratio of speedup to the number of processors - a linear speedup would correspond to a 100% efficiency. For 4 processors, the speedup is only 3.1 (an efficiency of 78%). However, for R80, the 4-processor efficiency is 86% (with a speedup of 3.4), and for R120, it is 90% (with a speedup of 3.6). The length of the time step was decreased for the higher resolution runs (we used 15 minutes for R40, 7.5 minutes for R80, and 5 minutes for R120). Finally, we reran the R80 timing tests on an 8-processor Cray Y-MP. The 8-processor run shows a speedup of 5.3 (an efficiency of 66%).

Table 4. Elapsed times for single processor (s.p.) and multiprocessor (m.p.) runs on the Cray 2 (C2) and Cray Y-MP (Y-MP). NCPUs is the number of processors of the m.p. run, speedup the ratio of s.p. to m.p. elapsed time, and efficiency the ratio of speedup to NCPUs.

Resolution	Computer	NCPUs	s.p. time	m.p. time	Speed up	Efficiency
R40	C2	2	135.84	72.77	1.9	93%
R40	C2	4	135.84	43.31	3.1	78%
R80	C2	4	1426.63	415.69	3.4	86%
R120	C2	4	6645.33	1852.70	3.6	90%
R80	Y-MP	8	727.10	138.00	5.3	66%

In interpreting the multiprocessing efficiencies, it is useful to consider the four sources of sub-linear speedup: incomplete parallelization of the computation, multiprocessing overhead, load imbalances among the different tasks, and memory contention problems. In the timings reported here, the overhead is negligible, as discussed above. The other sources of inefficiency are not easily isolated. Incomplete parallelization of the computation implies that there is a certain fraction of the code that is executed in serial, not parallel, mode. Load imbalances can occur if the computational work cannot be evenly divided among the processors. If all latitude tasks required the same amount of computations, the load imbalance could easily be estimated from the mismatch between the number of processors and number of latitude tasks. This kind of analysis was successfully performed for an adiabatic version of the GSM by Hoffman and Nehrkorn (1989). However, in the present case with full physics, the latitude tasks will require different amounts of CPU time, and load imbalances are not

easily determined without additional diagnostics. Memory contention problems can occur when two latitude tasks need to increment the same global accumulators. Again these could only be detected with the aid of special diagnostics. We have not implemented any diagnostics.

All the sources of inefficiency become more important as the number of processors increases relative to the problem size. Incomplete parallelization becomes more important as the number of processors is increased, because the parallel portion of the code requires a smaller fraction of the total wall-clock time. Load imbalances can become more severe, because as the size of the tasks assigned to the processors becomes smaller, the fractional differences between the loads increase. Memory contention becomes more likely as more and more processors try to access the same number of storage locations. This explains the clear trend in table 3 that for a given problem size, the efficiency decreases with increasing number of processors, and conversely for a given number of processors, the efficiency increases with increasing problem size.

We have indicated possible future enhancements in comments in the code, which might mitigate the memory contention problems. They would require calls to machine-specific parallel processing library routines, and should only be considered once it is determined that a particular platform will be available in dedicated mode for a large enough time to make it worth the extra effort.

3.3 Overall Speedups

To assess the overall effectiveness of our optimization effort, one has to combine the speedups due to improved single-processor execution speed with the speedups due to multiprocessing. Compared with the original code (at R40), the optimized, multiprocessed code exhibits a 5.9-fold speedup when run on 4-processor Cray 2. The computational cost of the enhanced physics evaluated in the context of the optimized code (a factor of 2.2 if one considers the different number of vertical levels) is more than made up by the multiprocessing speedup on even a 4-processor Cray 2 (a factor of 3.1). Thus, the goal of making the implementation of the enhanced physics package feasible was achieved by our optimization techniques.

Furthermore, since the R80 resolution 8-processor time on a Cray Y-MP is only 2% larger than the R40 single-processor time on the Cray 2, a doubling of the resolution would be possible at the same turn-around time if the computing platform were upgraded from a Cray 2 to a Cray Y-MP. The recently announced Cray C90 would present further opportunities for increases in resolution.

4. Forecast Results

We conducted several 4-day forecasts, using a FGGE 3b analysis for 12 UTC 12 January 1979 as the initial condition. This analysis was produced at ECMWF using a resolution of 1.875° . Boundary data sets were prepared using routines and databases supplied by PL. Forecasts were produced using the GSM with GWC physics, as described in Section 3, with a resolution of R40. The enhanced physics GSM (referred to as APGSM in the remainder of this section) was used for forecasts with a R40, R80, and R120 resolution. The size of the transform grid was determined, for each truncation, from the minimum requirements of exact transforms of quadratic terms, and, in the restrictions imposed by the FFTs (the number of longitudes must be a multiple of 2, 3, and 5). The resulting grid sizes are 128 longitudes by 102 latitudes for R40, 250 by 202 for R80, and 384 by 302 for R120. The vertical structure of the GSM was the same 18-layer structure for all forecasts, with sigma interfaces at 0, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.546, 0.642, 0.735, 0.820, 0.893, 0.948, 0.973, 0.990, and 1.0. The length of the time step was decreased in accordance with the increase in resolution (20 minutes was used for the R30 runs reported in I, and we used here 15 minutes for R40, 7.5 minutes for R80, and 5 minutes for R120). The other numerical constants that needed adjusting were the diffusion coefficients for the horizontal 4th order diffusion. At R30, a value of $6 \times 10^{15} \text{ m}^4/\text{s}$ was used in I. For the higher resolutions, we chose a value that would result in the same fractional damping per timestep of the highest wavenumber (approximately 6%). The corresponding values are $2.55 \times 10^{15} \text{ m}^4/\text{s}$ for R40, $3.23 \times 10^{14} \text{ m}^4/\text{s}$ for R80, and $9.60 \times 10^{12} \text{ m}^4/\text{s}$ for R120. The remaining numerical constants were left unchanged from one resolution to the next (for the timestepping, a time-filter coefficient of 0.04 and a time-step coefficient of 0.5, corresponding to a semi-implicit step, were used).

It should be noted that for a more realistic assessment of resolution on forecast skill, the physical parameterization would have to be retuned, since they all contain a number of adjustable parameters that have been tuned to the R30 resolution used most extensively by PL. In addition, the vertical structure should also be reconsidered, particularly in the case of the high resolution (R120) run. The results reported here should thus be considered a preliminary exploration of the effect of higher horizontal resolution.

Forecasts are verified against the FGGE 3b analyses in the following discussion. Global root mean square (rms) and mean (bias) errors for the GWC physics GSM are shown in Figure 1, at forecast lead times of 24, 48, 72, and 96 hours. Because we used the same 18-layer structure as in the APGSM forecasts, the drag-law scheme of the GWC GSM affected only a very thin bottom layer, leading to somewhat exaggerated forecast errors in the boundary layer (below $\sigma=0.9$). These large low-level errors are apparent in all panels of Figure 1. In the free atmosphere, the most striking feature of the error structure is the large positive temperature bias which grows with time (up to 3 K by day 4), presumably due to the lack of radiative cooling of the atmosphere in this model. This is also reflected in a sizeable geopotential height bias. Temperature rms errors also increase with forecast lead time to 4.8 K throughout most of the troposphere by day 4. The geopotential height rms error is dominated by the bias, reaching a maximum value of 137 m by day 4 at the tropopause (disregarding values for the top level of the model). Errors of the rotational wind exhibit maxima at the surface, at the top of the model, and at the jet level. At the jet level, the maximum values grow from 7.9 m/s to 19.9 m/s from day 1 to 4. Errors of the divergent wind are not shown, because the quality of the verifying analysis is uncertain.




Fig. 1: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) for the GWC GSM forecast with R40 resolution. Errors are shown at day 1 (solid line), day 2 (dotted line), day 3 (dashed line), and day 4 (dash-dotted line).





Fig. 1: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) for the GWC GSM forecast with R40 resolution. Errors are shown at day 1 (solid line), day 2 (dotted line), day 3 (dashed line), and day 4 (dash-dotted line).




Fig. 1: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) for the GWC GSM forecast with R40 resolution. Errors are shown at day 1 (solid line), day 2 (dotted line), day 3 (dashed line), and day 4 (dash-dotted line).

The corresponding figure for the APGSM R40 forecast (Figure 2) reveals sharp decreases in boundary layer errors, and the elimination of the positive temperature bias (which is now slightly negative in the upper troposphere, reaching -39 m after 4 days). The rms errors of temperature and height are decreased as well, to less than 4 K throughout most of the troposphere by day 4 in the case of temperature, and less than 80 m in the case of height. The rotational wind errors are smaller throughout the atmosphere, with jet level wind errors between 7.8 m/s (at day 1) and 16.5 m/s (at day 4).

For a direct comparison of day 4 forecast errors, the curves for the GWC GSM and the APGSM R40 forecasts are plotted together in Figure 3. Also shown in Figure 3 are the error curves for the APGSM R80 and R120 forecasts, allowing an assessment of physics and resolution changes. Errors of all variables are most affected by the change in physics packages, and only to a lesser degree by increases in resolution. The most obvious effect of the higher resolution is a reduction of the negative temperature and height bias in the upper troposphere. Rms errors of these quantities are also reduced overall by increases in resolution, but not consistently at all levels. Beneficial effects of increased resolution are not obvious for the rotational wind, which shows errors that are essentially identical for all three truncations.

It is not uncommon that increases in forecast skill, as measured by the kind of global statistics presented here, are hard to detect when model resolution is increased (Simmons, 1991, personal communication). The primary reason for this is that while low resolution forecasts will be missing small-scale features present in the verifying analyses (and the atmosphere), high resolution model forecasts may not predict them at the proper location (or time). In that kind of scenario, the high resolution forecast would be penalized more severely in terms of rms statistics. In addition, the verifying analysis used here is not high resolution, so small scales present in the R120, and, to some extent, the R80 forecasts would appear as errors. It is thus not surprising that in the tests reported here, where the high resolution model was also not properly tuned, the results are somewhat ambiguous.





Fig. 2: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) for the APGSM GSM forecast with R40 resolution. Errors are shown at day 1 (solid line), day 2 (dotted line), day 3 (dashed line), and day 4 (dash-dotted line).





Fig. 2: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) for the APGSM GSM forecast with R40 resolution. Errors are shown at day 1 (solid line), day 2 (dotted line), day 3 (dashed line), and day 4 (dash-dotted line).




Fig. 2: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) for the APGSM GSM forecast with R40 resolution. Errors are shown at day 1 (solid line), day 2 (dotted line), day 3 (dashed line), and day 4 (dash-dotted line).




Fig. 3: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) at day 4. Errors are shown for the GWC GSM forecast with R40 resolution (solid line), and the APGSM forecasts with R40 (dotted line), R80 (dashed line), and R120 (dash-dotted line) resolution.





Fig. 3: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) at day 4. Errors are shown for the GWC GSM forecast with R40 resolution (solid line), and the APGSM forecasts with R40 (dotted line), R80 (dashed line), and R120 (dash-dotted line) resolution.




Fig. 3: Global mean (a, c) and rms (b, d, e) error of temperature (a, b), geopotential height (c, d), and rotational wind (e) at day 4. Errors are shown for the GWC GSM forecast with R40 resolution (solid line), and the APGSM forecasts with R40 (dotted line), R80 (dashed line), and R120 (dash-dotted line) resolution.

5. Minimization of pressure gradient force errors

5.1 Pressure-gradient force errors

The low accuracy of geopotential calculations in spectral numerical weather prediction models introduces substantial errors, and therefore affects the overall forecast error. The order of magnitude of the errors commonly reaches 100 m at 500 mb level (in non-dissipative systems, Phillips, 1974), and generates erroneous systematic cooling of the atmosphere above the high mountains regions. This error is the most noticeable in conjunction with the calculations of *the Pressure-Gradient Force* (PGF)

$$- \nabla_{\sigma} \Phi - RT \nabla \ln p_s \quad (5.1.1)$$

which involves horizontal differencing of the geopotential, Φ (p_s is the surface pressure, R is universal gas constant, T is temperature and σ is the normalized pressure coordinate). Over sloping terrain the two terms of (5.1.1) tend to be large in absolute values and to have opposite signs. The cancellation, however, is difficult to achieve in the models, and a 1% error in temperature (2-3° C) will result in a 10% error in the pressure gradient force (Sundqvist, 1975). An example of the PGF existing at the model level 17 is given in Figure 4 (a) (see Section 5.3 for data description). This force is produced by the spectrally represented orography as demonstrated by the figure. The PGF is clearly shown over the high mountain regions; Himalayas, Andes and Greenland (max vector intensity plotted is 10^{-2} m/s^2). With an increase of the model spectral resolution, this PGF decreases as the mountains are represented better.

Together with the truncation error, the erroneous estimates of geopotential, Φ , are a persistent source of spurious PGF in the spectral models (Gary, 1973). The low accuracy of geopotential calculations is mainly caused by *the vertical differencing scheme* used to compute Φ from the model hydrostatic equation. If we integrate the hydrostatic equation in

the σ -coordinate system

$$\Phi(\sigma) = \Phi(1) - \int_1^\sigma RT d(\ln \sigma) \quad (5.1.2)$$

we see that the geopotential at any σ -level is only a function of temperature at and below that level. Therefore, the derivative $\partial\Phi / \partial(\ln\sigma)$ should not be approximated by a higher order finite-difference scheme in order to preserve the continuity of the hydrostatic equation (Janjic, 1977). As a result, the geopotential is calculated with a low accuracy that leads to erroneous estimates of (5.1.1). These inadequacies have been investigated in part by Brenner (1988), who showed that all commonly used schemes produced substantial height errors at the upper model levels. An example of the errors produced by the current model vertical differencing scheme (equation (15) in Brenner *et. al.*, 1982) is illustrated by Figure 4 (b). It shows the difference of the PGF intensity, at the model level 2, between the analytically determined and numerically determined Φ . In both cases the same spectrally represented orography and surface pressure are used. Again, the additional PGF, caused by the model vertical differencing scheme, exists in the high mountain regions. Note that, in this particular case, the difference of the PGF intensity is of order 1% (10^{-3} m/s^2) of the PGF in diagram (a).

5.2 Error minimization

These errors can be reduced if *the standard stratification approximation* is used, *i.e.*, if the thermodynamic model variable is not the temperature T , but the deviation of temperature T' , from some reference temperature T_r

$$T' = T - T_r \quad (5.2.1)$$

The reference temperature, T_r , can be defined to be a function that removes most of the cancellation in the calculations of term (5.1.1). It is determined





Figure 4. (a) Spurious pressure gradient force for a standard atmosphere due to spectral representation of mountains at R20 truncation. (b) Difference in spurious pressure gradient force at model level 2, due to model vertical differencing scheme for the same spectral truncation (R20).

by the stratification factor C_0 defined as

$$C_0^2 = \frac{R^2 T_r}{g} \left(\frac{g}{c_p} + \frac{dT_r}{dz} \right) \quad (5.2.2)$$

where, g is the gravity constant, c_p is the specific heat at constant pressure and z is the vertical coordinate. This factor can be chosen to be

$$C_0 = \begin{cases} \text{const.} \\ f(p) \\ f(p, \varphi) \end{cases}$$

i.e., constant, function of pressure only, or function of pressure and latitude φ (Zhang *et al.*, 1990). The model development at the ECMWF (Simmons and Chen, 1991) is based on defining the stratification factor as a function of vertical coordinate only, i.e., pressure or σ . The work of Zhang *et al.* includes a more general case of a reference atmosphere varying with latitude.

Two possible approaches in applying this approximation to the model have been investigated at the ECMWF. The first approach (called a TR scheme; Simmons and Chen [1991]) is based on the assumption that the reference atmosphere is removed from the spectrally represented thermodynamic variable. In this approach a substantial modification of the model code is required since the spectral variable is changed from the temperature itself to the perturbation of temperature, T' . A second approach, that requires less extensive modifications (identified as the PG-scheme), is based on using the standard stratification approximation in the model calculations that involve the PGF calculations only.

5.2.1 Schemes

In the present study we have tested three schemes: the PG-scheme originally used in the ECMWF study, a modified version of the PG-scheme called MPG-scheme, and the null scheme which consists of an isothermal reference profile. Table 5 summarizes the differences in parameter values for the three schemes.

Table 5 . Parameter values used in the three schemes: the PG, the MPG and null (an isothermal case).

Scheme	T_0 [°K]	p_0 [kPa]	Λ [°K/km]
PG	288	1013.2	6.5
MPG	270	1013.2	2.5
null	300	1013.2	0.0

5.2.1.1 The PG - scheme

This scheme is based on a reference atmosphere defined as

$$T_r(p) = T_0 \left(\frac{p}{p_0} \right)^\alpha \quad (5.2.3)$$

where T_0 and p_0 are the surface values of the temperature and pressure respectively, $\alpha = \Lambda R/g$ and Λ is the lapse rate. This reference atmosphere is subtracted from the model virtual temperature for the purpose of calculating the geopotential (see Section 5.2.2 for a description of the code modification).

We used this scheme in the present study since the idealized tests of the ECMWF model showed that the PG-scheme accounted for an error reduction similar to the reduction obtained by the more extensive TR scheme. The parameters were chosen to be the same as those reported in Simmons and Chen (1991) (the first row of table 5). In Figure 5, diagram (a) illustrates this reference temperature in respect to the standard atmospheric profile described in Section 5.3.

5.2.1.2 The MPG -scheme

We also considered, in the present study, a modified PG-scheme (the MPG -scheme) which is based on the same reference profile (5.2.3) but uses different parameter values (the second row of table 5). As Figure 5 (b) shows, this profile resembles a linear approximation of the standard temperature profile with an unrealistic lapse rate. The motivation for considering this modification was to estimate what impact the variation of reference atmosphere would have on the present model error. Further modifications of this scheme are not presented here since the improvements that we found were insignificant (see Section 5.4).

5.2.1.3 The null scheme - an isothermal case

Figure 5 (c) illustrates the case in which the reference atmosphere is isothermal. The value of the surface temperature T_0 was chosen to be 300°K which is the same as the reference temperature used in the model semi-implicit time integration scheme. This profile was used as a reference for an earlier version of the model when the thermodynamic variable was a perturbation from a uniform temperature profile.

5.2.2 Code modification

The implementation of the PG - scheme (5.2.3) requires a minimal modification of the code (Simmons and Chang, 1990). It involves only the equations that include PGF terms and, if applied to the grid-point thermodynamic variable, it is localized to a very few subroutines of the multi - tasking model code.




Figure 5. Reference temperature profiles compared to the temperature profile of the standard atmosphere, used in: (a) the PG scheme, (b) the MPG scheme and (c) the isothermal case.

5.2.2.1 Model equations

The model equation that directly involves PGF calculations is the divergence equation (Kaplan *et al.*, 1985)

$$\frac{\partial D}{\partial t} = -\frac{1}{a \cos^2 \phi} \left[\frac{\partial B}{\partial \lambda} - \cos \phi \frac{\partial A}{\partial \phi} \right] - \nabla^2 [E + \Phi] + \nabla \cdot F_h \quad (5.2.4)$$

where the notation is identical to Brenner *et al.* (1982). When the temperature variable, T , is replaced by the expression (5.2.1), the equation retains the same form for the temperature perturbation T' , with additional terms due to a reference atmosphere

$$- \left[\frac{1}{a^2 \cos^2 \phi} \frac{\partial}{\partial \lambda} R T_r \frac{\partial \ln p_*}{\partial \lambda} + \frac{1}{a^2 \cos \phi} \frac{\partial}{\partial \phi} R T_r \cos \phi \frac{\partial \ln p_*}{\partial \phi} \right] + \nabla^2 \int_1^\sigma c_p \kappa T_r d(\ln \sigma). \quad (5.2.5)$$

When compared to (5.1.1) we see that the later expression represents the PGF due to a reference atmosphere, T_r , in the model coordinate system. Furthermore, if T_r is replaced by the expression (5.2.3), the above expression reduces to one term

$$- \nabla^2 \frac{R T_0}{\alpha} \left(\frac{p_*}{p_0} \right)^\alpha. \quad (5.2.6)$$

This term can be easily combined with the geopotential calculations for the perturbation temperature under the Laplacian operator in equation (5.2.4). Therefore, the major cancellation of the PGF, due to a reference atmosphere, can be removed from the calculations in the divergence equation if the geopotential calculations in the hydrostatic equation are modified by the term (5.2.6). Indeed, this has been done by Simmons and Chen (1990), who modified the surface geopotential (equation

15) in the ECMWF hydrostatic equation by the term identical to the expression under the Laplacian operator in (5.2.6).

5.2.2.2 Code changes

For the experiments in this part of the study we used the dry-adiabatic version of the code. The present code structure enabled a relatively easy implementation of the changes described above. The changes were made in a grid-point space, which eliminated the need for additional corrections of the rest of the model equations and the semi-implicit time integration scheme. (In the case of implementing the schemes in the spectral space, it would be necessary to make quite radical model code changes.)

The first change that was made was to calculate the reference temperature in grid-point space. This was done in the subroutine NLPROD, that estimates nonlinear terms, where grid-point values of the surface pressure, p^* , are available. The reference temperature was then subtracted from the temperature field and added back to the field after the evaluation of the nonlinear terms.

The second and most significant change was to integrate the model hydrostatic equation in the grid-point space. (The current model formulation involves hydrostatic integrations in the spectral space.) Therefore, the subroutine consisting of the hydrostatic equation has been recoded and called before the nonlinear term evaluation. During this process, the surface geopotential was also modified by the corresponding portion of the term (5.2.6).

The third type of modification mostly involved the changes related to the postprocessing of model integrations. In particular, for each model run we have created evolution files for model tendencies, we have developed routines for calculations of the pressure gradient force, maximal grid-point tendencies and spectral modes that have the maximal amplitude at each model level (this enabled us to check if any of the schemes exhibit spectral biases). Finally, we have modified the postprocessing package to plot the

forecast fields at the model σ -levels instead of the mandatory pressure levels. The latter was needed in order to eliminate further contribution of the hydrostatic equation to the temperature (and geopotential) errors in the post-processing cycle.

5.3 Initial Data

Most of the results presented here are from idealized tests. As in Simmons and Chen, idealized tests were based on initial data describing a dry atmosphere at rest. They consisted of an analytically prescribed temperature profile which is a function of pressure alone, and which is of a form such that the balanced surface pressure could be computed analytically given the basic model orography, Φ_s . In our study, we used the analytical profiles defined by Phillips (1974)

$$T\left(\ln \frac{p}{p_0}\right) = \frac{1110}{R} \left\{ 72.43 - \ln \frac{p}{p_0} \left[2(-6.9) - 3 \ln \frac{p}{p_0} \right] \right\} \quad (5.3.1)$$

An example of this temperature profile is shown in Figure 5. This temperature profile has been extensively used in numerical studies related to spurious PGF. The balanced surface pressure, p_s , was computed from

$$\Phi_s\left(\ln \frac{p_s}{p_0}\right) = 1110 \left\{ 0.95 - \ln \frac{p_s}{p_0} \left[72.43 - \ln \frac{p_s}{p_0} \left(-6.9 - \ln \frac{p_s}{p_0} \right) \right] \right\} \quad (5.3.2)$$

To calculate the surface pressure we used the surface geopotential from the gridded FGGE 3b data set. The orography data was smoothed horizontally first, and then truncated spectrally at the needed resolution.

5.4 Tests

The idealized tests were created with the purpose of investigating how well the PGF-scheme minimizes present model error. Simmons and Chang found that the errors were significantly reduced, especially for longer time integrations, at truncation T43. The degree of forecast improvement was such that the accuracy of T43 model runs became equivalent to the unmodified T106 model runs. These are significant corrections that can affect the cost of model integrations substantially. The purpose of this study is to determine if the similar improvements of the forecasts can be achieved by the present model at R40 truncation and compared to the forecasts with R80 truncation.

Before starting the long, ten-days model runs at R40 and R80 truncations, we tested the minimization schemes at a lower spectral truncation, R20 and R40. This was based on the fact that the PGF errors are so tightly related to the spectral truncation error (as described in Section 5.1.) that the effects of the schemes should be noticeable even at the lower spectral resolution. In general, we produced four types of model runs. The Control runs were based on an unmodified model code and used as a reference for the "PG" and the "MPG" runs, which utilized the PG and the MPG schemes, respectively. Finally, "Isothermal" runs were based on an isothermal reference atmosphere. We also tested the schemes' performance in respect to distribution of the model vertical levels (equally-spaced levels *vs.* current vertical structure) and the model diffusion. The results are presented in terms of the model initial tendencies and the one-day and ten-days forecasts.

We limited the investigation to the cases described above, and we reviewed the results of the idealized runs only. We did not find the substantial improvements of the forecasts by the minimization schemes considered, although the model vorticity dynamics showed some sensitivity to the presence/absence of the reference atmosphere.

5.4.1 Scheme comparison

The performance of the three schemes are best illustrated by the initial tendencies presented in Figure 6. Two diagrams on the left-hand side show max grid-point divergence tendency at R20 (upper panel) and R40 (lower panel) model truncation. It is noticeable that the PG and MPG schemes do not differ very much from the unmodified code, while the "isothermal" scheme produces the largest initial tendencies. Two panels on the right-hand side of the figure show the corresponding vorticity tendencies. Although the differences between the schemes are not large, some sensitivity of the vorticity is noticeable. In particular, for the R20 case the "isothermal" scheme produces the smallest tendencies, while in the R40 case (the lower panel), the PG-scheme produces the smallest tendencies. Also, in the case of the PG-scheme, it appears that the maximum tendency decreases slightly with height (with exception of the model level 15) in R20 truncation, but increases slightly with height (except at the model level 13) in R40 truncation.

Since the differences in Figure 6 appear to be very small, we considered the most sensitive, the velocity fields, for the forecast comparison. Figure 7 shows the differences of the global RMS of the velocity field. Again, the upper two panels represent R20 truncation, and the lower two panels represent R40 truncation. The differences between the forecasts are of order of 0.1 cm/s. In the case of the divergent velocity component (two panels on the left-hand side of the figure) the PG and the MPG produce larger values than the control (solid line patterns) at R20 truncation. At R40 truncation, both of the schemes produce larger values at the upper model levels (except the PG-scheme at the level 3), and slightly smaller values at the lower model levels. The rotational wind component (two panels on the right-hand side of the figure) shows that both of the schemes produce larger values at all model levels except the uppermost levels (5-1 and 3-1, respectively).

The geographical distribution of the winds at level $\sigma = 0.124$ for R20 forecasts is shown in the Figure 8. Diagram (a) shows a distribution of the wind speed obtained in the Control case after 10 days of integrations. The




Figure 6. Vertical distribution of maximum initial tendencies for two model truncations: R20 (upper panels) and R40 (lower panels). Divergence tendencies are shown on the left-hand side and vorticity on the right-hand side.




Figure 7. Vertical distribution of the velocity RMS. The upper two panels apply to R20 truncation, and the lower two panels apply to R40 truncation. Differences between the PG, MPG and Control runs are plotted. Two panels on the left-hand side are derived for the divergent component of the velocity, while panels on the right hand-side are derived for the rotational part of the velocity.

wind maxima are noticeable over the Himalayas, Greenland and the Antarctica. The lower diagram (b) shows the difference in wind speed between the Control forecast and the PG-scheme (the PG-scheme produced smaller wind speeds).

Figure 9 shows the winds for R40 truncation at $\sigma = 0.594$. Diagram (a) shows the horizontal wind field after 1-day of integrations in the Control case (maximum wind does not exceed 5m/s). The rotational component of the wind is pronounced over the mountains. Again, diagram (b) shows the difference in the wind speed between the Control case and the PG-scheme. It is noticeable from the last two figures that the differences are very small and localized above the mountains.

Comparison of the global mean values of the differences between forecasted and initial temperature and geopotential fields corresponding to Figure 7 is shown in Figure 10. In general, the 10-days forecasts at R20 produced slightly increased temperature (of the order of 10^{-2}) at the upper model levels ($\sigma = 0.497$ to $\sigma = 0.021$) and decreased temperature at the lower levels ($\sigma = 0.995$ to $\sigma = .594$). In these forecasts, the schemes tended to produce smaller increase at the upper levels, and almost no difference in the forecasts at the lower levels (diagrams on the left of Figure 10). This is reflected in difference between the geopotential forecasts to be of the order of 10^{-2} m. Similar behavior was obtained at R40 truncation for one day forecasts (two lower panels).

The comparison of the vertical distribution of the spectral modes, that have the largest amplitudes, did not show any particular bias of the schemes. This is illustrated by table 6.




Figure 9. Geographical distribution of the horizontal wind for R40 model truncation. Diagram (a) shows the winds at model level 11 after 1-day integration in the Control case. Diagram (b) shows the difference between the wind speed forecasts obtained with the PG-scheme and unmodified code.




Figure 10. Vertical distribution of global mean values of the differences between initial data and forecasts. Two panels on the left-hand side are derived for the temperature, while panels on the right hand-side are derived for the geopotential. The upper two panels apply to 10-days forecast R20 truncation, and the lower two panels apply to 1-day forecasts at R40 truncation. Differences between the PG, MPG and Control runs are plotted. (Note difference between the horizontal axes of the upper and the lower panels).

Table 6. Vertical distribution of the vorticity spectral wave number (m,n) that has max amplitude after one day of integrations.

Model truncation	R20			R40		
Scheme	Control	PG	MPG	Control	PG	MPG
σ						
.021	(0,12)	(0,12)		(0,7)	(0,7)	(0,7)
.074	(1,16)	(1,16)		(1,12)	(1,12)	(1,12)
.124	(3,21)	(3,21)		(0,39)	(0,39)	(0,39)
.174	(1,14)	(1,14)		(3,21)	(3,21)	(3,21)
.225	(1,20)	(1,20)		(8,35)	(8,35)	(8,35)
.275	(1,20)	(1,20)		(7,35)	(7,35)	(7,35)
.325	(7,25)	(7,25)		(0,17)	(0,17)	(0,17)
.375	(7,25)	(7,25)		(2,20)	(2,20)	(2,20)
.425	(8,25)	(8,25)		(0,18)	(0,18)	(0,18)
.497	(8,25)	(8,25)		(0,18)	(0,18)	(0,18)
.594	(0,5)	(0,5)		(0,29)	(0,29)	(0,29)
.688	(3,17)	(3,17)		(0,29)	(0,29)	(0,29)
.777	(3,17)	(3,17)		(0,29)	(0,29)	(0,29)
.856	(10,28)	(10,28)		(8,31)	(8,31)	(8,31)
.920	(10,28)	(10,28)		(3,17)	(3,17)	(3,17)
.960	(10,28)	(10,28)		(3,17)	(3,17)	(3,17)
.981	(10,28)	(10,28)		(0,20)	(0,20)	(0,20)
.995	(10,28)	(10,28)		(0,20)	(0,20)	(0,20)

5.4.2 Vertical distribution of the model levels

As indicated by the results reviewed in the previous section, some variability of the schemes with the height exists. To examine the impact of the model level distribution on error minimization schemes, we repeated the Control, the PG and the MPG runs with equally spaced model levels at R20 truncation. The results are shown in Figure 11.




Figure 11. Vertical distribution of maximum initial tendencies for the model truncation R20, and different distribution of the model levels. The results shown on the left-hand side are obtained with the current distribution, and the results shown on the right-hand side are obtained with equally-spaced model levels.

Two panels on the left-hand side of the figure are equivalent to the upper two panels of Figure 6. They represent the differences between the maximum grid-point initial tendencies obtained with the current distribution of model vertical levels. The other two panels, on the right-hand side of the figure, are corresponding tendencies that were obtained with the equally-spaced model levels. It is noticeable, from the figure, that the performance of the minimization schemes is not affected by the level distribution in respect to the control runs. Slight differences are noticeable in the vorticity field (the lower two panels), but not of a significant value.

5.4.3 The model diffusion

Since the differences reviewed in earlier two sections were small, we compared them with the amount of horizontal diffusion applied to the Laplacian term in the equation (5.2.4) that involves geopotential estimates. The results of this comparison are given in Figure 12.

An evolution of total divergence and vorticity tendencies during 1-day forecasts at R20 truncation are plotted on the diagrams. Diagrams on the left-hand side of the figure represent the Control and the PG runs without diffusion, while the diagrams on the right-hand side represent the corresponding runs with applied horizontal diffusion. The upper two diagrams represent the divergence tendencies, and the lower two diagrams represent the vorticity tendencies. It is quite noticeable that the scheme follows the Control run quite closely independently on the model horizontal diffusion. Differences appear to be slightly larger for the divergence tendencies than for the vorticity tendencies. The horizontal axis represents the number of time steps 30 min long.




Figure 12. Evolution diagrams of total divergence tendencies (the upper diagrams) and total vorticity tendencies (the lower diagrams) from the runs with and without horizontal diffusion. Model truncation is R20, and the length of forecasts is one day.

5.5 Conclusions

In the present study we investigated schemes for minimization of the PGF errors in the current version of the global spectral model. We tested three schemes, named the PG-scheme, the MPG-scheme and the Iso-thermal case against the Control case that consisted of the unmodified model code. The forecasts were made at the R20 and R40 model truncation and were of durations varying from 1-hour to 1-day to 10 days. The initial data described a dry standard atmosphere that justified use of dry adiabatic model code without any physics.

We tested the error minimization against the various scheme parameters (the temperature and the lapse rate), the model spectral resolution (R20 and R40), the distribution of the model vertical levels (current *vs.* equally-spaced levels) and the model horizontal diffusion.

The idealized tests showed that no substantial improvements of the forecasts were achieved by the PG-scheme. Some sensitivity was found in the fields related to the model vorticity dynamics (the rotational velocity and the vorticity). The changes found did not exceed 1% of the Control forecasts obtained by the current model code. We did not find this justifiable for extending the study to the integrations at higher spectral resolution (*i.e.*, R80). The results obtained at the lower resolution are qualitatively in agreement with the results of idealized tests of Simmons and Chen (1990). The only difference is that the ECMWF orography is more realistic and spectrally fitted to T106, while the orography available for the runs with the current model version has been smoothed by the low spectral truncation. This resulted in smaller initial tendencies than the ones obtained by the ECMWF.

Although substantial improvements of the forecasts by the PG-scheme were not found, the integrations with the real data and more realistic orography may provide improved forecasts. Simmons and Chen have found improvements in real data forecasts which were described later by specific vorticity dynamics of the ECMWF model (personal communication with Simmons). In our idealized tests, it was found that the additional

PGF force has been created by the vertical differencing scheme at the same geographical positions where the ECMWF forecast improvements were located. Together with the sensitivity of the rotational wind and the vorticity field, which we also diagnosed, this can be indicative of possible forecast improvements with the real data and orography, similar to the ones obtained in the ECMWF study.

6. Summary

The development of a vectorized, multiprocessing global spectral model (GSM) with enhanced physical parameterizations is documented. The starting point was the Phillips Laboratory GSM, in its GL90 version, containing an enhanced suite of physical parameterizations. The latitude tasking scheme for multiprocessing the loop over latitude in the calculation of the spectral tendencies and adjusted model variables was implemented, using the general truncation version of the hydrodynamics code. Wave-number calculations were vectorized over wavenumber and multiprocessed over vertical level. All gridpoint calculations were vectorized over longitude, and the physics packages were brought into closer compliance with plug-compatibility rules.

Speedups due to the optimization were demonstrated in single- and multiprocessing timing tests on a dedicated Cray 2 and Cray Y-MP. Compared with the original code (at R40), the optimized code runs 1.9 times faster on a single processor Cray 2. When run on 4-processor Cray 2, the optimized code exhibits a 5.9-fold speedup compared to the original version. The computational cost of the enhanced physics evaluated in the context of the optimized code (a factor of 2.2 if one considers the different number of vertical levels) is more than made up by the multiprocessing speedup on even a 4-processor Cray 2 (a factor of 3.1). Thus, the goal of making the implementation of the enhanced physics package feasible was achieved by our optimization techniques.

Furthermore, since the R80 resolution 8-processor time on a Cray Y-MP is only 2% larger than the R40 single-processor time on the Cray 2, a doubling of the resolution would be possible at the same turn-around time if the computing platform were upgraded from a Cray 2 to a Cray Y-MP. The recently announced Cray C90 would present further opportunities for increases in resolution.

The advanced physics GSM (APGSM) was evaluated in forecast tests, and contrasted with the simpler physics GSM used at GWC. Use of the APGSM leads to sharp decreases in boundary layer forecast errors, and the

elimination of the positive temperature bias of the GWC GSM (which is now slightly negative in the upper troposphere, reaching -39 m after 4 days). The rms errors of temperature and height are decreased as well, and the rotational wind errors are smaller throughout the atmosphere. Forecast errors of all variables are most affected by the change in physics packages, and only to a lesser degree by increases in resolution. The most obvious effect of the higher resolution is a reduction of the negative temperature and height bias in the upper troposphere. Rms errors of these quantities are also reduced overall by increases in resolution, but not consistently at all levels. Beneficial effects of increased resolution are not obvious for the rotational wind, which shows errors that are essentially identical for all three truncations.

Minimization of errors in the computation of the horizontal pressure gradient force was investigated, using a perturbation temperature instead of full temperature in the integration of the hydrostatic equation. Several schemes were tested in an idealized model atmosphere. The idealized tests showed no substantial improvements of the forecasts were achieved by any of the schemes we tested. Some sensitivity was found in the fields related to the model vorticity dynamics (the rotational velocity and the vorticity). However, the changes found did not exceed 1% of the control forecasts obtained by the current model code.

7. References

- Brenner, S., 1988: A comparison of various numerical solutions of the hydrostatic equation. Proc. 8th Conference on Numerical Weather Prediction, Baltimore, MD, American Meteorological Society, Boston, MA, 438-441.
- Brenner, S., C.-H. Yang and K. Mitchell, 1984: The AFGL global spectral model: Expanded resolution baseline version. Report No. **84-0308**. AFGL, Hanscom AFB, MA 01731. 72 pp. [NTIS ADA160370].
- Brenner, S., C.-H. Yang and S. Y. K. Yee, 1982: The AFGL spectral model of the moist global atmosphere: Documentation of the baseline version. Report No. **82-0393**. AFGL, Hanscom AFB, MA 01731. 65 pp. [NTIS ADA129283].

- Hoffman, R. N., 1985: An implementation of the semi-implicit time scheme for primitive equation spectral models. *Mon. Wea. Rev.*, **113**, 1829-1831.
- Hoffman, R.N., and T. Nehr Korn, 1989: Multiprocessing a Global Spectral Numerical Weather Prediction Model, in Proceedings of the Fourth SIAM Conference on Parallel Processing for Scientific Computing, Dec. 11-13, 1989, Chicago, IL., 168-173.
- Kalnay, E., M. Kanamitsu, J. Pfaendtner, J. Sela, M. Suarez, J. Stackpole, J. Tuccillo, L. Umscheid, and D. Williamson. Rules for interchange of physical parameterizations. *Bulletin of the American Meteorological Society*, **70**:620-622, 1989.
- Kaplan, L.D., J.-F. Louis, R.N. Hoffman, R.G. Isaacs, W.J. Gutowski, and T. Nehr Korn, 1985: Improving numerical weather prediction by maximizing the use of assimilated satellite data. Tech. Rep. **85-0298**, Air Force Geophysics Laboratory, Hanscom AFB, MA. [NTISADA169295].
- Liou, K. N. and G.D. Wittman. Parameterization of the radiative properties of clouds. *Journal of the Atmospheric Sciences*, **36**:1261-1273, 1979.
- Liou, K. N., S. C. Ou, S. Kinne, and G. Keonig. Radiation parameterization program for use in general circulation models. Technical Report **84-0217**, Air Force Geophysics Laboratory, Hanscom AFB, MA, 1984. [NTIS ADA148015].
- Mahrt, L., H.-L. Pan, J. Paumier and I. Troen, 1984: A boundary layer parameterization for a general circulation model. Report No. **84-0063**. AFGL, Hanscom AFB, MA 01731. 179 pp. ADA144224
- Mahrt, L., H.-L. Pan, P. Ruscher and C.-T. Chu, 1987: Boundary layer parameterization for a global spectral model. Report No. **87-0246**. AFGL, Hanscom AFB, MA 01731. 188 pp. [NTIS ADA199440]
- Nehr Korn, T., R.N. Hoffman, and J.-F. Louis, 1990: Design of an enhanced global spectral model. Tech. Rep. **90-0309**, Geophysics Laboratory, Hanscom AFB, MA. ADA232123
- Norquist, D., and C.-H. Yang, 1990: Refinement and testing of the moist convection parameterization in the GL global spectral model. Tech. Rep. **90-0285**, Geophysics Laboratory, Hanscom AFB, MA. [NTIS ADA 241684]

- Ou, S.-C. and K.-N. Liou, 1988: Development of radiation and cloud parameterization programs for AFGL Global Models. Report No. 88-0018. AFGL, Hanscom AFB, MA 01731, 88 pp. [NTIS ADA193369] .
- Phillips, N.A., 1974: Application of Arakawa's Energy Conserving Layer Model to Operational Numerical Weather Prediction., Office Note 104, National Meteorological Center, NWS/NOAA, 40 pp.
- Sela, J. G., 1980: Spectral modeling at the National Meteorological Center. *Mon. Wea. Rev.*, **108**, 1279-1292.
- Simmons, A. J., and J. Chen, 1990: The calculation of geopotential and the pressure gradient in the ECMWF atmospheric model; influence on the simulation of the polar atmosphere and on temperature analyses. ECMWF Technical Report No. 66, Shinfield Park, Reading, pp. 76.
- Simmons, A. J., and J. Chen, 1991: The calculation of geopotential and the pressure gradient in the ECMWF atmospheric model: influence on the simulation of the polar atmosphere and on temperature analyses. *Q. J. R. Meteorol. Soc.*, **117**, 29-58.
- Slingo, J. M., 1987: The development and verification of a cloud prediction scheme for the ECMWF model. *Quart. J. R. Meteor. Soc.*, **113**, 899-928.
- Soong, S.-T., Y. Ogura and W.-S. Kau, 1985: A study of cumulus parameterization in a global circulation model. Report No. 85-0160. AFGL, Hanscom AFB, MA 01731. 113 pp. [NTIS ADA170137].
- Tiedtke, M., 1989: A comprehensive mass flux scheme for cumulus parameterization in large-scale models. *Mon. Wea. Rev.*, **117**, 1779-1800.
- WMO, Geneva. Report of the Seminar on Progress in Numerical Modelling and the Understanding of Predictability as a Result of the Global Weather Experiment, Sigtuna, Sweden, October 1984. TD 33. [Publication date 1985].
- Zhang, D., H. Seng and L. Ji, 1990: Development and testing of hydrostatic extraction scheme in spectral model. *Advances in Atmos. Sci.*, **7**, 142-153.

Appendix A: User's Guide

The document reproduced on the following pages is included as a plain text file (named "users.guide") in the software deliverables of this contract.

Users's Guide to the Global Spectral Model (GSM)

1. Overview

This document describes the Global Spectral Model (GSM), in its general truncation version, as configured for the GL (new) physics with vectorization and multiprocessing. A technical report (GL-TR-90-0309, NTIS ADA232123) describes the algorithms and software design of this version. For a scientific documentation, refer to the GL technical reports, and the references contained therein (hydrodynamics: AFGL-TR-82-0393, NTIS ADA129283; AFGL-TR-84-0308, NTIS ADA160370; PBL: AFGL-TR-84-0063, NTIS ADA144224; AFGL-TR-87-0246, NTIS ADA199440; PL-TR-91-2031, NTIS ADA235310; radiation: AFGL-TR-84-0217, NTIS ADA148015; AFGL-TR-88-0018, NTIS ADA193369; convection: AFGL-TR-85-0160, NTIS ADA170137; GL-TR-90-0285, NTIS ADA241684). This code employs what is known as the "GL90" physics with the following changes. The 3-decker radiation scheme is used in version 5.1; version 5.2 uses an improved radiation scheme ("rad5") along with the tuned Slingo parameterization of stratiform and convective cloud cover (Slingo, 1987: QJRM, 113, 889-927). This document applies to both version 5.1 and version 5.2; differences between the two versions exist in the boundary data set format, and are noted in the appropriate section. Minor changes and bug fixes have been made to pbl and Kuo schemes. In the following, the installation and running of the GSM are discussed.

2. Portability

The code itself is strictly FORTRAN 77 (ANSI X3.9-1978), with the following exceptions:

- INCLUDE statements are used throughout the code to include common blocks and parameter statements; the syntax of the statements (and filenames) may differ from one machine to the next, but (almost) all machines support this Fortran extension
- Multiprocessing (refer to C<mp> comment lines) is implemented using CRAY autotasking directives (CMIC\$ comment lines). These are comments in standard fortran. In addition the Cray extension task common is used to identify private common blocks. In a standard single processor version of the code these would be ordinary common blocks.
- Module tcheck.f uses machine-dependent system routines for obtaining the time and date, and CPU time
- Some of the modules contain edit symbols that need to be replaced before compilation. (Edit symbols are strings enclosed by <> brackets.)

Along with the code several files are provided for compiling, linking, and running the code. These are developed for a UNICOS operating system, and are easily adapted to different machines using UNIX. They are:

makefile - used by the UNIX make facility for compilation and linking of the gsm (executable stored in gsm.x) and the utility program asize (stored in asize.x). The replacement of the edit symbols is included in the compilation rule where appropriate - the UNIX sed facility is used for this purpose.

Changes needed for porting: the directory pathnames, and the compiler statements and flags.

gsm.sed - sed script used by the makefile.

Changes needed for porting: for 32-bit machines, edit symbols <DOUBLE>, <DP> should be set for double precision, for 64-bit machines they should be set for single precision.

rungsm - c-shell script for running the gsm

Changes needed for porting: directory pathnames

3. Configuration

3.1. Arrangement of code

The source code for the GSM is stored in a number of different files. The fortran code is stored in files with extension ".f". Modules containing edit symbols are stored in files with extension ".fed"; edited versions of these modules (with the edit symbols replaced by running the sed script gsm.sed) are stored in files with extension ".f". The fortran modules contain include statements for the inclusion of common and parameter statements to facilitate consistent use of commons and easy changes to array dimensions. Include blocks are stored in files with extension ".blk". Include blocks containing edit symbols are stored in files with extension ".bed"; after replacement of edit symbols, the edited include blocks are stored in files with extension ".blk". A complete list of the file names is appended to this document. The makefile also contains the current list of source and object file names.

Aside from the GSM code proper, the FFTs are loaded in separately as a library. Since these may be more readily changed from one installation to the next, they are maintained in a separate directory (and with a separate makefile). They do not contain edit symbols, and do not use any of the GSM include blocks.

Edit symbols in the code (strings enclosed by <>) are used to allow two options that cannot be accommodated with standard fortran:

- the model can be run in a hemispheric or a global version; this requires use of the edit symbols "<GLO>" and "<HEM>", which serve to activate lines of code. For the global version, "<GLO>" is replaced by five blank spaces, and "<HEM>" by "C", and vice versa for the hemispheric version.

- certain sensitive calculations can be performed in single or double precision. Generally, single precision should be used on 64-bit machines, and double precision on 32-bit machines. To select single precision, the symbols "<DOUBLE>" and "<DP>" are replaced by "REAL" and "E"; for double precision, they are replaced by "DOUBLE PRECISION" and "D".

The replacement of the edit symbols is handled automatically by the makefile (using the sed script gsm.sed). To change any one of the above two options, this script must be modified accordingly. The editing step may be eliminated if the options are not expected to be changed - changes in resolution will not require edit symbol replacements.

3.2. Resolution changes

The program allows great flexibility in the choice of spectral truncation and vertical resolution. Changing either requires changing parameter values in the include blocks containing the parameters (gsmdim.blk and param.blk), in the work space include block (work0.blk), and, if applicable, changing the data statement for the vertical levels (in block data vrtst0, in gdata.f). In addition, there are some vertical resolution dependent constants in the radiation scheme which must be changed in the initialization entry point of that module.

Although it is relatively painless to change the resolution in the GSM we caution the reader that the parameterizations have been tuned for rhomboidal 40 truncation, using 18 layers as described by Norquist (GL-TR-90-0285). Increasing resolution may require retuning or possibly even reformulating the physics. Otherwise it is possible that increasing resolution will degrade forecast skill. In addition constants describing the time step and diffusion coefficients should be changed to be compatible with the new resolution. Finally the input boundary value data sets will need to be recreated at the new resolution. These aspect of changing the model resolution are not addressed in this document.

Any pentagonal spectral truncation can be selected by choosing the following three parameters (in param.blk):

NMAX	the maximum value of the meridional index n
MMAX	the maximum value of the zonal wavenumber m
NMAX0	the maximum value of n at m=0

The popular rhomboidal (R) and triangular (T) truncations are selected by letting

rhomboidal (e.g., R30): NMAX0=MMAX (=30), NMAX=NMAX0+MMAX (=60)
 triangular (e.g., T30): NMAX0=MMAX=NMAX (=30).

In addition, the following two parameters allow the computational domain to be a subregion of the globe

PER	the periodicity in the zonal direction (=1 for a global domain)
-----	---

NHEM =1 for hemispheric, =2 for the global version

For a global domain (covering both hemispheres and all longitudes), PER=1 and NHEM=2. The value of NHEM must be consistent with the choice of the <HEM> and <GLO> edit symbols in gsm.sed.

The vertical resolution is controlled by the two parameters (in gsmdim.blk):

KPDIM number of layers
KHDIM number of layers which carry humidity.

The arrangement of the levels in the vertical is set by a data statement in block data vrtst0, in gdata.f, for the array

SI(H) sigma at level H.

This data statement must be consistent with a top-down vertical structure, i.e. SI(1)=0 and SI(KP+1)=1.

Based on the selection of these seven resolution-determining parameters (NMAX through KHDIM), several other parameters related to the number of spectral coefficients and the size of the transform grid must be determined. A standalone program (asize.f) exists to compute values for remaining parameters that need to be specified in gsmdim.blk, param.blk and work0.blk. The variables described here, as well as all other important variables in the GSM, are all described in the comments of the file containing the common blocks in which they are stored. Further a complete global list of these variables is presented in the appendix to this guide.

3.3. Input and output files

The input and output to and from the GSM involves files of name 'tapenn', where nn is the unit number (the only exception to this rule is output to unit IOTERM: no file name is specified for this unit, allowing the use of "standard out" if the appropriate value for IOTERM is used). The c-shell script (rungsm) provides the link between these internal file names and user-chosen file names. The unit numbers are stored in cntrl.blk, and initialized in cntrl0 (in gdata.f). The following lists the variable name along with the current value of the input and output units:

INPUT:

IOCNTL=99: control input - read in by CTRLID
IOGSM1=11: GSM initial conditions spectral coefficients input data set (either for initial start - read in by RDINIT, or for a restart - read in by RHIST)
IOGSM7=17 GSM surface data input data set (not needed for restart) - read in by RDINIT

OUTPUT:

IOTERM=6: (std out) printed output for console
IOLIST=16: printed output for list file
IOGSM2=12: GSM forecast spectral coefficients output data set - written out by WDATA, or WDATA2 (if only one time step and not an NMI run), or WNMI (if an NMI run)

IOGS11=21: GSM restart output data set (written by WHIST)

In addition, there are the following internal files, and diagnostic output files (some of which are not currently written to):

IOGS15=15: GSM holding file for radiation
IOGS25=25: GSM diagnostic output file for radiation
IOGS26=26: GSM diagnostic output file for Kuo convection
IOGS27=27: GSM diagnostic output file for Kuo convection
IOGS50=50: GSM diagnostic output file for pbl.

The file structure of the input and output files is described in more detail in the following. The file structure conforms with current GL practice, with the following exceptions:

- control input:

Istart greater than or less than zero indicates a restart or NMI calculation.

Ntime less than or equal to zero indicates preliminary values or tendencies from the first time step will be output.

Floating point values are used for the record 2 variables (trms, hfinc, hrsinc) described below. If hrsinc is less than zero, the history restart output file is appended to rather than rewind each time a restart record is output. This option was absent from the GL GSM.

- spectral coefficient data set:

The precipitation record described below is written out entirely whenever the spectral arrays are written.

- history restart data set:

The header record contains the value of nstep.

This data set may be appended to (see description of hrsinc).

3.3.1. Control input

Read in by CTLRID (in gmain.fed)

The control input file allows setting a number of values for the current GSM without having to recompile it. They are read in using list-directed I/O from an ASCII file, in the following order:

Record 1:

istart: restart/nmi run flag (>0 for restart, <0 for nmi run, =0 for normal initial start)

ntime: no. of time steps for this run (integer)
ntime = or < 0 implies that tendencies or preliminary values will be written after the first, unfiltered time step, either by wnmi (if istart<0) or wdata2 (if istart=0)
ntime= 0, -1, -2, -3 implies that STEP1 will return after the

```

        call to laloop, diffsn, implct, conpcp, respectively.
        idt: time step length (integer secs)
Record 2:
        trms: rms summary print interval (real hrs)
        HFINC: fcst spectral coefficient output interval (real hrs)
        HRSINC: restart history file output interval (real hrs)
                (by default, the history file will be rewound before each
                output - irewnd=1 -; use hfinc=-abs(hfinc) for irewnd=0)
Record 3:
        mf: input moisture adjustment flag ( 1 for yes, 0 for no )
        fv: moisture adjustment coefficient (real - e.g. 0.9)
Record 4:
        andree: time filter coefficient (real - e.g. 0.04)
        alpha: semi-implicit parameter (real - between 0.5 and 1.0)
Record 5:
        dkh: horizontal diffusion coefficient for divergence (real)
        ekh: horizontal diffusion coefficient for other variables (real)

```

To do a restart, set `istart=number of filtered time steps of the restart data set`, and `ntime=number of timesteps left that are necessary to finish the forecast`:

```
ntime(restart fcst) = ntime(total fcst) - istart.
```

3.3.2. Spectral coefficient data set

Read in by RDINIT; written out by WDATA, WDATA2.

This is a binary data file, consisting of the following:

```

Record 1: (Header record)
Four variables: nstep (integer), htime (real),
                itime (character*8), idate (character*8)
They have the following meaning:
nstep - number of filtered time steps completed (a negative
        number is used to denote an unfiltered data set)
htime - forecast hour of data set
itime(1:4) - initial state time in 'hh2' format
idate(1:8) - initial state date in 'ddmonyy' format (where mon is
        the 3-letter, lower case month designator; other formats
        require changes to the tempus subroutine in gmain.fed)

```

Initial data sets (as read by RDINIT) must have `htime=0` and `nstep=0`.

At present, the only time an unfiltered data set is written out in this format is by WDATA2, which is called if `ntime = or < 0` and `istart=0` is specified in the input (see the discussion under control input): depending on the value of `ntime`, all or part of the first time step is completed, and either tendencies or unfiltered values are written out. In any case, `nstep` in the file is -1, and `htime=dt`.

Records 2 - 6: (spectral coefficients of data)

The spectral coefficients of absolute vorticity, divergence, temperature, and specific humidity (all at KP layers), and of

$\ln(\text{surface pressure}/1000 \text{ mb})$, corresponding to *nstep* and *htime*.

NOTE:

For each layer, all spectral coefficients are present for a rhomboidal truncation with $m \leq MMAX$, $(n-m) \leq NMAX0$, with *m* increasing along the diagonals. On input, the inactive spectral coefficients are ignored; on output, they are output as zeros.

Record 7: (surface geopotential)

The spectral coefficients of the surface geopotential.

Record 8: (Precip - only present for *nstep* > 0)

This record contains 3 global arrays (on the transform grid - *nlont***nhem***nlatt* values) plus 4 global accumulators (*nstep*+1 values each, one for each time step) - all are floating point arrays:

- geprev - total precip at time *htime*
- gecpv - convective precip at time *htime*
- evpmp - scaled evaporation at current time step
- geptot, cvptot - total, convective global precip at each time step
- nkuo - number of gridpoints with Kuo convection at each time step
- evptot - total evaporation at each time step.

NOTE:

To avoid memory contention problems, the actual dimension and number of locations used in many arrays are two separate parameters. For example, the number of longitudes of the transform grid (NLONT) may be 128. On the CRAY2 using NLONT to dimension arrays would be disastrous in this case. Therefore the actual array dimension (NLONTD) is taken to be 129. However, file storage is always in terms of the actual number (NLONT). Therefore, depending on the values chosen for NLONT and NLONTD, the internal storage of the gridpoint arrays may differ from the storage in file - this is handled automatically by the I/O routines. Thus, the value of NLONTD (but not of NLONT) may be changed without having to modify the input files.

3.3.3. NMI data set

Written out by WNMI.

This is a binary data file, consisting of the following:

Record 1: (Header record)

Four variables: *nstep* (integer), *htime* (real),
 itime (character*8), *idate* (character*8)

They have the following meaning:

- nstep* - number of filtered time steps completed (a negative number is used to denote an unfiltered data set)
- htime* - forecast hour of data set
- itime*(1:4) - initial state time in 'hhZ' format
- idate*(1:8) - initial state date in 'ddmonyy' format (where mon is the 3-letter, lower case month designator; other formats require changes to the tempus subroutine in gmain.fed)

At present, the only time a data set is written out in this format is by WNMI, which is called if ntime = or < 0 and istart < 0 is specified in the input (see the discussion under control input): depending on the value of ntime, all or part of the first time step is completed, and either tendencies or unfiltered values are written out. In any case, nstep in the file is -1, and htime=dt.

Records 2 - 6: (spectral coefficients of initial data)

These records are in the same format as records 2-6 of the spectral coefficient data set, except that the values of the initial data are written out (valid at htime=0).

Records 7 - 11: (spectral coefficients of tendencies or unfiltered values)

These records are in the same format as records 2-6 of the spectral coefficient data set, except that the values of the tendencies or unfiltered values are written out (valid at htime=dt).

Record 12: (surface geopotential)

The spectral coefficients of the surface geopotential.

3.3.4. Surface data set

Read in by RDINIT

This is a binary data file, consisting of the following:

Record 1: (Header record)

Four variables: nstep (integer), htime (real),
itime (character*8), idate (character*8)

These are not used inside the program (other than being echoed to the list file), and are used only for descriptive purposes.

for version 5.1:

Records 2 - 10: (global fields of data)

for version 5.2:

Records 2 - 5 and 9 - 13: (global fields of data)

These records contain values of the following fields on the transform grid (nlont*nhem*nlatt values), at 1 or 2 levels (nlev=1 or 2):

nlev

1	ts - surface (skin) temperature
1	cdr - drag coefficient
1	salb - solar albedo
1	esd - equivalent snow depth
2	stc - soil temperature
1	ws - surface moisture
2	smc - soil moisture
1	canopy - vegetation cover

1 stbot- temperature below lowest soil layer
for version 5.2 only:
Records 6 - 8: Fixed fields used by radiation
These are zonal mean fields, at nhem*nlatt latitudes and npmrad levels
(npmrad = 41 at present) (ordered as j=1,nhem*nlatt, k=1, npmrad):

clpr - pressure levels for the next 2 fields
clh2o- climatological values of water vapor
clo3 - climatological values of Ozone

3.3.5. Restart data set

Read in by RHIST; written out by WHIST.

This is a binary data file, consisting of the following:

Record 1: (Header record)

Four variables: nstep (integer), htimel (real),
itime (character*8), idate (character*8)

They have the following meaning:

nstep - number of filtered time steps completed (this must
be a positive number)

htimel - forecast hour of the filtered time step

itime(1:4) - initial state time in 'hh2' format

idate(1:8) - initial state date in 'ddmonyy' format (where mon is
the 3-letter, lower case month designator; other formats
require changes to the tempus subroutine in gamin.fed)

for version 5.1:

Records 2 - 10: Identical to records 2-10 of the surface data set

for version 5.2:

Records 2 - 13: Identical to records 2-13 of the surface data set

Records 14 - 16: Global fields of convection parameters used by
radiation

The global fields are written as (i=1,nlont, j=1,nhem*nlatt)

cnvpre - convective precipitation rate

kcnvtop- layer index of convective cloud top

kcnvbot- layer index of convective cloud bottom

The record numbers in the following are correct for version 5.1. For
version 5.2, they have to be incremented by 6.

Records 11 - 10+(nlatt+1)/2: (radiative heating rates)

These (nlatt+1)/2 records contain the radiative heating rates and the
surface flux. Each record contains nlont*nhem*(kp+1) values. These
data are necessary for a restart at a time step without a full
radiation calculation. The internal storage of the data differs from
the storage in the input/output data set. The file format was chosen
for compatibility with the previous version of the GL GSM.

NOTE: in the original GL version, there were only nlatt/2 records

Records 11+(nlatt+1)/2 - 15+(nlatt+1)/2: (spectral coefficient data at

time htime1)

These records are in the same format as records 2-6 of the spectral coefficient data set.

Record 16+(nlatt+1)/2: (surface geopotential)

The spectral coefficients of the surface geopotential as record 7 of the spectral coefficient data set.

Record 17+(nlatt+1)/2: (Precip valid at time htime1)

This record contains 5 global arrays (on the transform grid - nlont*nhem*nlatt values):

- geprev - total precip at time htime1
- gecpv - convective precip at time htime
- evmp - scaled evaporation for current time step
- dewmp - accumulated total dew
- etotbl - accumulated total evaporation

Records 18+(nlatt+1)/2 - 22+(nlatt+1)/2: (spectral coefficient data at time htime2)

These records are in the same format as records 2-6 of the spectral coefficient data set, except they contain data for the unfiltered time step (valid time htime2=htime+dt, time step number nstep+1)

Record 23+(nlatt+1)/2: (Precip valid at time htime2 plus global accumulators)

This record contains 2 global arrays (on the transform grid - nlont*nhem*nlatt values) plus 4 global accumulators (nstep+1 values each, one for each time step) - all but nkuo are floating point arrays:

- geshem - total precip at time htime2
- gecnv - convective precip at time htime2
- geptot,cvptot - total, convective global precip at each time step
- nkuo - number of gridpoints with Kuo convection at each time step
 - ** integer array - note difference from spectral coefficient data set **
- evptot - total evaporation at each time step.

3.4. Multiprocessing

The code is designed to be multiprocessed on a shared memory machine. The GSM contains comments of the form "C<mp> ... " in several modules. These comments are meant to help in the implementation of the multiprocessing, and provide inline documentation after implementation.

We have implemented multiprocessing as Cray microtasking directives, and task common statements. The CRAY autotasking directives are activated by invoking cf77 with the -Zp option. Generally speaking autotasking automatically makes maximum use of available resources. However the UNICOS command setenv NCPUS may be used to restrict the GSM to fewer than the full set of available processors.

All common blocks in the GSM are global and should be shared among the processors, except for the following, which are local to each task and are identified in the code as task common. These task common blocks contain data for a single latitude.

```

/aleg/, in aleg.bed - Legendre functions at a single latitude
/grdpt/, in grdpt.blk - gridpoint values of variables at one latitude pair
/prod/, in prod.blk - gridpoint values of nonlinear terms at one
                    latitude pair
/work1/,/worksp/, in work0.blk - work space common
/pblval/, in pblvar.blk - pbl variables local to each latitude
/abci/, in gpbl.f - pbl variables local to each latitude

```

Results with multiprocessing are not strictly reproducible because the order of addition in the Gaussian sums (i.e. the grid point to spectral transforms) is arbitrary. Differences between runs due to this are initially the size of round off error, but may amplify exponentially. Runs using a single processor are reproducible. In addition, for testing purposes only, we have included some code to enforce the order of the sums. This code, named spin.f, is commented out and is extremely inefficient. We have ourselves used it only to make sure that the differences between multiprocessing runs is in fact due to differences in the order of the sums and not some other problem.

We note that the radiation parameterization performs a full calculation only for odd latitudes. The results for an even latitude are simply taken to be that calculated for the neighboring, equatorward odd latitude. Consequently the multiprocessing for the main latitude loop in laloop, parcels out the latitudes in pairs. Changes to the radiation parameterization which increase or decrease the coupling between latitudes will generally have implications for the multiprocessing strategy. For example if the radiation calculation was performed in full for each latitude, then the multiprocessing in laloop could be done latitude by latitude singly.

4. List of filenames

The following is a complete list of the file names of the GSM code modules currently in use:

4.1. Include blocks with edit symbols:

```

aleg.bed - Legendre functions at a single latitude
bl.bed - surface geopotential
cpcnst.bed - computational and physical constants
dtime.bed - spectral coefficients of time level 3
ipmn.bed - constants used for Pmns calculation
time1.bed - spectral coefficients of time level 1
time2.bed - spectral coefficients of time level 2

```

4.2. Include blocks without edit symbols

```

char.blk - character constants
clock.blk - variables related to time stepping and valid time

```

cntrl.blk - control parameters
 ctauw.blk - global array used by modkuo
 czalb.blk - albedo and internal radiation array
 dercot.blk - common used by dercot.f
 fftn0.blk - constants used by FFTs
 grddim.blk - parameter statements used in radiation package
 grdpt.blk - gridpoint values of variables at one latitude pair
 ground.blk - global arrays of surface parameters
 moisgl.blk - precipitation values
 gsmdim.blk - PARAMETER statements, included in param.blk
 param.blk - PARAMETER statements
 pblvar.blk - commons used in pbl package
 prod.blk - gridpoint values of nonlinear terms at one latitude pair
 ridiaq.blk - commons used in radiation package
 tconst.blk - constants for time scheme
 vrtst.blk - vertical structure information
 waven.blk - wavenumber arrays
 wavvec.blk - arrays needed for gtranv (vectorizing Gaussian sums)
 work0.blk - work space common

For version 5.2:

mstrad.blk - convection parameters used by radiation

4.3. Fortran modules with edit symbols

gconvec.fed - Module referenced by gmdkuo.f (Kuo convection)
 gcphys.fed - Cover routine for adjustment physics
 ginit.fed - Initialization routines for hydrodynamics
 gmain.fed - Main routine and time step driver and I/O routines
 gtend.fed - tendency calculation (hydrodynamics) routines
 gtrans.fed - version of gtranv.fed for completely general truncation
 (Gaussian sums will not vectorize)
 gtranv.fed - spectral transform routines; includes cover routine
 for FFTS, but not the FFTs themselves
 gtstep.fed - time stepping routines
 gutil.fed - utility subroutines

4.4. Fortran modules without edit symbols

dercot.f - Derickson and Cotton code for saturation vapor pressure
 dumphy.f - dummy routine for convection
 dumrad.f - dummy routine for radiation
 gbldum.f - dummy routine for PBL
 gdata.f - block data statements
 gxtet.f - dry adiabatic adjustment
 glrgsc.f - large scale precipitation
 gmdkuo.f - Kuo convection
 gpbl.f - pbl package
 grad.f - radiation package
 gradia.f - radiation package
 rdsdum.f - dummy version of rdsfc-gl.f
 rdsfc-gl.f - module to read in boundary data for pbl and radiation
 spin.f - module to enforce same order of Gaussian sums
 svp.f - module for computing saturation vapor pressure
 tcheck.f - checkpointing and error handling routines
 work0.f - work space allocation routine

4.5. Other related files

gsm.sed - sed script for replacing edit symbols
makefile - makefile for compiling and linking GSM
rungsm - c-shell script for running gsm
asize.f - standalone program for computing dimension parameters

0. Define outline-mode:

```
;;; Local Variables: ***  
;;; mode:outline ***  
;;; outline-regexp:"\\([0-9]+\\.\\|\\|)+" ***  
;;; End: ***
```

Appendix B: List of Variables

The document reproduced on the following pages is included as a plain text file (named "variables.list") in the software deliverables of this contract.

Variable dictionary

Alphabetical list of commons, include file names (without .blk), variables in commons, and parameters. Also shown are the include file names where they are defined, and a short description. Descriptions that apply to several variables are only shown for the first one in this list - others are cross-referenced. Additional descriptions can be found in the include files and the fortran source code itself. This list is applicable to version 5.1.

Variable	Include file	Description
a	prod	$a = Z*U + \dots$ Eq. (3)
ae	param	radius of the earth.
aleg	aleg	aleg contains values of the associated Legendre functions, their meridional derivatives, etc., and the constants needed in the quadrature at a single latitude. All variables are defined at a single Northern Hemisphere latitude.
alfa	pblvar	alpha constant which determines if time scheme is backwards implicit (alpha=1) or semi-implicit (alpha=0.5)
alpha	tconst	
amtrx	tconst	A matrix Eq. (B4)
andre	tconst	beta: time filter factor
b	prod	$b = Z*V + \dots$ Eq. (4)
beta	pblvar	time filter factor
bl	bl	D matrix appearing in Eq. (B11) as the form $A*C + R*T0*(transpose DEL)$ where DEL appears in the text as r.
bmtrx	tconst	
bowen	pblvar	canopy - canopy moisture ybar(L,IHEM) the vertical integral from sigma = 0 to 1 at longitude L in hemisphere IHEM of Y = C and D where C is the advection of Q, i.e. velocity*gradient(Q).
canopy	ground	
cbar	prod	
char	char	contains character variables
clock	clock	
cmtrx	tconst	contains various parameters related to the forecast valid time C matrix defined following Eq. (B3)
cntrl	cntrl	
		contains variables defining forecast time parameters, and the I/O units, including,
cp	param	the specific heat of dry air at constant pressure
cpblnm	pblvar	contains computational and physical constants and all the values of the weights, GWGT, and the sines of the quadrature latitudes, GLAT. The transform grid (GLAT,GWGT) is always Gaussian.
cpcnst	cpcnst	
ctauw	ctauw	contains a grid point common

cth1	vrtst	block of horizontal convergence. It is computed in nlprod, and used in modkuo cth1(K),CTH2(K) constants for computing level temperatures used in Eq. (39).
cth2	vrtst	
cvptot	moisgl	
czalb	czalb	contains gridpoint arrays of surface parameters needed by the PBL routine
d	grdpt	x(L,K,IHEM) the value at longitude L in hemisphere IHEM in layer K of X = Z, D, T, W, U, and V (W defined only for moisture-carrying layers).
d1b	time1	x1b(IB,K) the IBth spectral coefficient in layer K of X = Z, D, T and W. (n-m=even)
d1c	time1	x1c(IC,K) the ICth spectral coefficient in layer K of X = Z, D, T and W. (n-m=odd)
d2b	time2	x2b(IB,K) the IBth spectral coefficient in layer K of X = Z, D, T and W. (n-m=even)
d2c	time2	x2c(IC,K) the ICth spectral coefficient in layer K of X = Z, D, T and W. (n-m=odd)
daccff	param	For direct access file, factor by which the number of (real) words in a record must be multiplied to obtain the record length.
dbar	prod	see under cbar
dddtb	dtime	dxdtb(IB,K) the IBth spectral coefficient tendency in layer K of X= Z, D, T and W. (n-m=even)
dddtc	dtime	dxdtc(IC,K) the ICth spectral coefficient tendency in layer K of X= Z, D, T and W. (n-m=odd)
del	vrtst	del(K) thickness of layer K Fig. (1)
delt	tconst	DEL(K)=SI(K+1)-SI(K)
delta	tconst	difference in time
dercot	dercot	DELTA*ALPHA
dew	pblvar	Common block for the saturation vapor pressure calculation
dewmp	moisgl	
dkh	cpcnst	accumulated total dew
dqdl	grdpt	diffusion coefficient for D
dqdp	grdpt	dqdy(L,IHEM) the value at longitude L in hemisphere IHEM of the derivative of Q w.r.t. Y = L and P which are lamda and phi, i. e. the longitude and latitude respectively:
dqdt	pblvar	cos(phi)*gradient(Q)=(DQDL,DQDP)
dqdtb	dtime	see under dqdl
		dqdtb(IB) the IBth spectral coefficient tendency of Q (n-m=even)

dqdtc	dtime	dqdtc(IC) the ICth spectral coefficient tendency of Q (n-m=odd)
dsoil	pblvar	
dt	tconst	time step (s)
dtddt	pblvar	
dtddtb	dtime	see under dddtb
dtddtc	dtime	see under dddtc
dtime	dtime	contains either the time tendencies of the prognostic variables or at the end of the time step the variables themselves at time level 3 (this is the time level following the central time):
dtimdb	dtime	
dtsoil	pblvar	
dudt	pblvar	
dvdt	pblvar	
dwtdb	dtime	see under dddtb
dwtdc	dtime	see under dddtc
dwsodt	pblvar	
dzdtb	dtime	see under dddtb
dzdtc	dtime	see under dddtc
e	prod	e = Kinetic energy/unit mass Eq. (5)
ekh	cpcnst	diffusion coefficient for Z, T and W.
ep	pblvar	
eps1	ipmn	
eps1on	ipmn	
eps1r	cpcnst	eps*latent heat of vaporization/R This is the constant a as defined following Eq. (77).
esatdc	dercot	
escl	cpcnst	eps*b
esc2	cpcnst	(eps - 1)*b eps is 0.622, the ratio of the molecular weights of water vapor and dry air.
esdc	dercot	
etotbl	moisgl	accumulated total evaporation
evap	pblvar	
evpmp	moisgl	scaled evaporation for current time step
evptot	moisgl	not used, except in I/O; initialized to zero)
fdown	pblvar	
fftn0	fftn0	
freqrd	param	contains the constant parts of the FFTs. frequency with which full radiation is computed (in hours)
gdc	dercot	
gecnv	moisgl	convective accumulated precip at time htime2
gecvp	moisgl	convective precip at time htime1
geprev	moisgl	total precip at time htime1
geptot	moisgl	geptot,cvptot - total, convective
geshem	moisgl	global precip at each time step
		total accumulated precip at time htime2

gflx	pblvar	
glat	cpcnst	glat(LA) sine of Gaussian latitudes (transform grid)
gninv	tconst	stored matrix inverses needed by semi-implicit time scheme. (See comments in IMPLCT.)
grav	param	g, acceleration of gravity
grddim	grddim	Defines several parameters for the dimensions of arrays and for i/o units. The important thing is that the radiation calculation is performed on a separate grid with ip longitudes and jp latitudes.
grdpt	grdpt	contains grid point values for a single Gaussian latitude of the variables:
ground	ground	contains gridpoint arrays of surface and ground parameters needed by the PBL routine. These arrays are read in at the start of, and carried along during the forecast.
gs2rd	cpcnst	gs2rd, GS7E5 boundary layer flux constants: $GS2RD = \sqrt{2} * (g * \sigma(KP)) / (R * DEL(KP))$ $GS7E5 = 7E-5 * 2 * (g * \sigma(KP)) / (R * DEL(KP))$ where KP is the layer next to the surface.
gs7e5	cpcnst	see under gs2rd
gsmdim	gsmdim	Parameters of array dimensions used in plug-compatible modules, and in param.blk:
gwtg	cpcnst	gwtg(LA) Gaussian quadrature weights (transform grid)
hab	aleg	hab(IB) $ILAPB(IB) * HB/AE$
hac	aleg	hac(IC) $ILAPC(IC) * HC/AE$ PAB and HAB are used to calculate UMS and VMA, PAC and HAC are used to calculate UMA and VMS, in Eqs. (32) and (33).
hb	aleg	hb(IB) the value of the meridional derivative of PB: $HB = -\cos(\phi) * d(PB) / d(\phi)$.
hc	aleg	hc(IC) the value of the meridional derivative of PC: $HC = -\cos(\phi) * d(PC) / d(\phi)$.
hdc	dercot	
hdf	pblvar	
heat	pblvar	
hfinc	cntrl	Forecast output increment
hflen	cntrl	Forecast length (total)
hfzero	cntrl	=1 to output initial conditions on IOGSM2
ho	clock	ho, idy, mo, iyr - hour (hh), day (dd), month (mm), and year (yy) of forecast initialization time (GMT)
hpbl	pblvar	

hpinc	cntrl	Precipitation output increment (NOT USED)
hpzero	cntrl	Precipitation accumulator zeroing increment
hrsinc	cntrl	Restart output increment
htime1	cntrl	Forecast time of variables in /TIME1/
htime2	cntrl	Forecast time of variables in /TIME2/
ib00	cpcnst	value of index IB for which m=0 and n=0
ib0n	waven	array of array indices IB for which m=MB(IB)=0
ic01	cpcnst	value of index IC for which m=0 and n=1
ic0n	waven	array of array indices IC for which m=MC(IC)=0
icl	ridiag	
idate	char	idate,ITIME Initial state date/time
idum1	cntrl	idum9 - IDUM1 Dummy entries for additional I/O units, this allows units to be added without changing the common block size
idum2	cntrl	see under idum1
idum3	cntrl	see under idum1
idum4	cntrl	see under idum1
idum5	cntrl	see under idum1
idum6	cntrl	see under idum1
idum7	cntrl	see under idum1
idum8	cntrl	see under idum1
idum9	cntrl	see under idum1
idy	clock	see under ho
iendw0	work0	last element of work array allocated
ifax	fftn0	
iheat	ridiag	
iheati	ridiag	
iheats	ridiag	
ilapb	waven	ilapb(IB) inverse Laplacian operator for n=NB(IB)
ilapc	waven	ilapc(IC) inverse Laplacian operator for n=NC(IC)
iocntl	param	Control input
iogs11	cntrl	GSM restart output data set
iogs15	cntrl	GSM holding file for radiation
iogs25	cntrl	GSM holding file for radiation
iogs26	cntrl	GSM diagnostic output file for Kuo convection
iogs27	cntrl	GSM diagnostic output file for Kuo convection
iogs50	cntrl	GSM diagnostic output file for pbl
iogsm0	cntrl	GSM tendencies s. c.
iogsm1	cntrl	output data set (NOT USED)
		GSM initial conditions spectral coefficients input data set (either for

iogsm2	cntrl	initial start or for restart)
		GSM forecast spectral
		coefficients output data set
iogsm7	cntrl	GSM surface data input data set
iogsm8	cntrl	GSM precipitation output
		data set (NOT USED)
iolist	param	List output
ioterm	param	Terminal output
ipdim	grddim	
ipmn	ipmn	contains the constant parts of the
		PMNS computation. Refer to routines
		INIPMN and PMNS (in GTRANS) for more
		details.
isfcf	ridiag	
isoil	pblvar	
istrw0	work0	istrw0(I) first element of WORK0
		allocated at subroutine level I
itldc	dercot	
it2dc	dercot	
itime	char	see under idate
iuerr	grddim	
iuout	grddim	
iy	clock	see under ho
jpblp	pblvar	
jpdim	grddim	
julday	clock	corresponding julian day
kflux	ridiag	
kh	param	number of layers which carry
		humidity (=khdim from gsmdim.blk)
kh1	param	lower index of humidity loop, i. e.
		there is moisture in layers
		K=KH1, ..., KP
khdim	gsmdim	kh: no of moist sigma layers
kp	param	number of layers (=kpdim
		from gsmdim.blk)
kpdim	grddim	kp: no of sigma layers
kpdim	gsmdim	
lapb	waven	lapb(IB) Laplacian operator
		corresponding to NB(IB)
lapc	waven	lapc(IC) Laplacian operator
		corresponding to NC(IC)
latpbl	pblvar	
lbdiag	wavvec	lbdiag(IBON) - Number of even spectral
		coefficients along the diagonal
		n-m=2*(IBON-1)
lcdiag	wavvec	lcdiag(ICON) - Number of odd spectral
		coefficients along the diagonal
		n-m=2*(ICON-1) + 1
lonpbl	pblvar	
maxstp	moisgl	
mb	waven	mb(IB) zonal wavenumber m corresponding
		to the IBth even spectral coefficient
mbldim	gsmdim	maximum no of sigma layers
		used in pbl
mc	waven	mc(IC) zonal wavenumber m corresponding
		to the ICth odd spectral coefficient

mmax	param	the maximum value of m
mmb	waven	mmb(IB) index in the Fourier coefficient arrays that corresponds to the wavenumber m=MB(IB)
mmc	waven	mmc(IC) index in the Fourier coefficient arrays that corresponds to the wavenumber m=MC(IC)
mo	clock	see under ho
moisgl	moisgl	contains gridpoint arrays of precip etc, and precip counters
moist	moisgl	
mpdim	grddim	nlontd: longitude dimension of transform grid arrays
mzbl	pblvar	
nb	waven	nb(IB) meridional index n corresponding to the IBth even spectral coefficient
nb0n	param	length of IB0N array (in /WAVEN/)
nc	waven	nc(IC) meridional index n corresponding to the ICth odd spectral coefficient
nc0n	param	length of IC0N array (in /WAVEN/)
ndiff	cpcnst	minimum n wavenumber to which diffusion operator is applied for variables other than D
ndimi0	work0	size of array ISTRW0
ndimw0	work0	size of work array
nhdim	gsmdim	nhem: no of hemispheres
nhem	param	1 for hemispheric, =2 for the global version (=nhdim from gsmdim.blk)
nkdim	grddim	
nkuo	moisgl	number of gridpoints with Kuo convection at each time step
nlatt	param	number of Gaussian latitudes in one hemisphere of transform grid, equator and pole never included. (=npdim/nhem, npdim from gsmdim.blk) (must be $\geq (2*NMAX0+NMAX+1)/4$ for MMAX $\leq 2*(NMAX-NMAX0)$ and $\geq (3*NMAX+1)/4$ for MMAX $\geq 2*(NMAX-NMAX0)$)
nlev0	work0	last element of ISTRW0 used, number of subroutine levels in use
nlon0	fftn0	
nlont	param	number of longitudes in transform grid. (must be $\geq 3*NUMZW-2$)
nlontd	param	longitude dimension of transform grid arrays (=mpdim from gsmdim.blk) (must be $\geq NLONT$)
nmax	param	the maximum value of n
nmax0	param	the maximum value of n at m=0
npbdim	pblvar	
npbl	pblvar	
npdim	grddim	nhem*nlatt: no of latitudes of transform grid
nrstp	cntrl	Time step to begin forecasting
nsodim	gsmdim	no of soil layers used in pbl (nsoil) + 1
nsoil	pblvar	
nstep	clock	time step number; during STCPN, the

unfiltered forecast is complete for NSTEP steps, and the unfiltered forecast for step NSTEP+1 (NT) is being computed; after STEP_N, the filtered forecast for nstep is complete (and stored in time level 1), and the unfiltered one for NSTEP+1 is complete (and stored in time level 2). The corresponding forecast hours are stored in /CNTRL/ as HTIME1 and HTIME2. (During STEP1 nstep=0; there is no filtering on the initial time (HTIME1=0), but the data is treated the same way as a filtered forecast after STEP1)

nt	clock	NSTEP + 1
numb	param	number of spectral coefficients with n-m=even
numbd	param	dimension for arrays of s. c. with n-m=even (must be >=NUMB)
numbs	param	maximum of NUMBD, NUMCD
numc	param	number of spectral coefficients with n-m=odd
numcd	param	dimension for arrays of s. c. with n-m=odd (must be >=NUMC)
numcs	param	maximum of NUMBD, NUMCD
numzw	param	number of zonal wavenumbers: wavenumber zero is included in this count. (=MMAX/PER + 1)
omega	param	
pab	aleg	pab(IB) ILAPB(IB)*PB/AE
pac	aleg	pac(IC) ILAPC(IC)*PC/AE
param	param	contains PARAMETER statements for the general truncation version of the GSM.
pb	aleg	pb(IB) the value of the IBth associated Legendre function for which n-m=even
pblbic	pblvar	
pblbig	pblvar	
pblk	pblvar	
pblvag	pblvar	
pblval	pblvar	
pblvar	pblvar	
pc	aleg	pc(IC) the value of the ICth associated Legendre function for which n-m=odd
per	param	the periodicity in the zonal direction
phisb	bl	
phisc	bl	
precip	pblvar	
prod	prod	contains the grid point values at a single latitude of the nonlinear products:
psbl	grdpt	
psfc	grdpt	psfc(l,ihem) surface pressure (in mb)
	pblvar	
qlb	timel	qlb(IB) the IBth spectral coefficient of Q. (n-m=even)
qlc	timel	qlc(IC) the ICth spectral coefficient of Q. (n-m=odd)

q2b	time2	q2b(1B) the 1Bth spectral coefficient of Q. (n-m=even)
q2c	time2	q2c(1C) the 1Cth spectral coefficient of Q. (n-m=odd)
qml	pblvar	
qs	ground	surface moisture
radintd	ridiag	
rdate	char	rdate, RTIME Run date/time (time forecast run on computer)
rex1	vrtst	rex1(K), REX2(K) ratios of Exner function in layer K to that in layers K-1 and K+1 respectively. These constants are used in the dry adiabatic adjustment.
rex2	vrtst	see under rex1
rgas	param	R, gas constant for dry air Virtual temperature effects are not included in this model.
rh	pblvar	
ri	prod	Non-flux part of thermodynamic equation. Eq. (29) (RI in the code is I in Brenner et al. (1982).)
ridiag	ridiag	
rkappa	cpcnst	kappa, R/Cp
rt0	tconst	$R \cdot T_0$ (m/s)**2
rtcs	cpcnst	cosine(alpha)
rtime	char	see under rdate
rtsn	cpcnst	sine(alpha) (alpha is the rotation angle between the surface wind and the wind in the lowest layer.)
salb	czalb	solar albedo
si	vrtst	si(H) sigma at level H
sikap	vrtst	sikap(H) sigma**kappa at level H
sinlat	aleg	the sine of the quadrature latitude in the Northern Hemisphere: GLAT(LA) or ALAT(LA)
sl	vrtst	sl(K) sigma in layer K
slkap	vrtst	Eq. (34) slkap(K) sigma**kappa in layer K
snoflx	pblvar	Eq. (34)
snow	ground	equivalent snow depth
ss	pblvar	
t	grdpt	see under d
t0	vrtst	t0(K) T0 in layer K
t1b	timel	see under d1b
t1c	timel	see under d1c
t2b	time2	see under d2b
t2c	time2	see under d2c
taox	pblvar	
taoy	pblvar	
tauw	ctauw	
tbl	grdpt	xbl(l,k,ihem) the value at timel of x(t), w, u, v (actual, not pseudo-velocities)
tbot	ground	temperature below lowest soil layer

tconst	tconst	contains constants needed for the implicit time scheme.
time1	time1	contains prognostic variables at time level 1 (the the time level preceeding the central time):
time2	time2	contains prognostic variables at time level 2 (the the central time level):
tml	pblvar	
trigs	fftn0	
ts	ground	surface (skin) temperature
tsoflx	pblvar	
tsoil	ground	soil temperature
tsoml	pblvar	
u	grdpt	see under d
ubl	grdpt	see under tbl
uml	pblvar	
ut	prod	ut = U*T : eastward temperature flux. Eq. (16)
uw	prod	uw = U*W : eastward moisture flux. Eq. (17)
v	grdpt	see under d
vbl	grdpt	see under tbl
vml	pblvar	
vertst	vertst	contains information about the vertical structure and constants used in vertical advection and the dry adiabatic adjustment.
vt	prod	vt = V*T : northward temperature flux. Eq. (16)
vw	prod	vw = V*W : northward moisture flux. Eq. (17)
w	grdpt	see under d
w1b	time1	see under d1b
w1c	time1	see under d1c
w2b	time2	see under d2b
w2c	time2	see under d2c
waven	waven	contains the zonal wavenumber m, meridional index n, and related quantities.
wavvec	wavvec	This common block is used by the vectorizing version of the transform routines (gtranv.fed). It assumes that the spectral truncation is pentagonal, and all spectral coefficients within the pentagon and with the given periodicity PER are active (i.e., the "arbitrary truncation" of AER TM NWP-1 is not supported). Furthermore, the following routines and this common block assume that the spectral coefficients are stored with m increasing along the diagonals defined by $n-m=i$, where i goes from 0 to $2*(NB0N-1)$ (even coeff.) or 1 to $2*(NC0N-1)+1$ (odd coeff.): GQB, GQC, LSUMB, LSUMC, ADVB, ADVC.

wbl	grdpt	see under tbl
wdf	pblvar	
wgtlat	aleg	the quadrature weight: GWGT(LA) or AWGT(LA)
wj	prod	wj = non-flux transport of moisture. Eq. (30) (WJ in the code is J in Brenner et al. (1982).)
wkt	pblvar	
work0	work0	contain the work array and related variables.
work0	work0	work array
work1	work0	
worksp	work0	contain the work array and related variables.
wsoflx	pblvar	
wsoil	ground	soil moisture
wsoml	pblvar	
x1	pblvar	
z	grdpt	see under d
z0	ground	roughness length
z1b	time1	see under d1b
z1c	time1	see under d1c
z2b	time2	see under d2b
z2c	time2	see under d2c
zlev	pblvar	