

U.S. DEPARTMENT OF COMMERCE
NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION
NATIONAL WEATHER SERVICE
NATIONAL CENTERS FOR ENVIRONMENTAL PREDICTION

OFFICE NOTE 388

G R I B

(Edition 1)

THE WMO FORMAT
FOR
THE STORAGE OF WEATHER PRODUCT INFORMATION
AND
THE EXCHANGE OF WEATHER PRODUCT MESSAGES
IN GRIDDED BINARY FORM
AS USED BY
NCEP CENTRAL OPERATIONS

Clifford H. Dey
NCEP Central Operations

Revised
(see overleaf)

March 10, 1998

This is an unreviewed manuscript,
primarily intended for informal exchange of information
among NCEP staff members

NCEP Office Note 388 -- GRIB

REVISION HISTORY

All revisions in the text, of any substance, are marked by a vertical bar to the left of the correction location, in the same manner as this paragraph. Insertions, deletions, and alterations are all so marked. Deleted text will not appear, of course, in the printed pages, but the location of the deleted material is marked by the sinister bar.

Please insert (or replace) the indicated pages in your copy of this document.

Date	Section	Page	Nature of change
Dec. 14, 1993	0	5	Clarification of Indicator Section content
Dec. 21, 1993	A	5	Correction of forecast hour typos
Dec. 22, 1993	4	4	Clarification of count of second order values
Feb. 18, 1994	1	4,5	Added generating processes 5 and 88
	1	12	Corrected table of 1.25x1.25 points per row
	2	13	Corrected summation limit
Mar. 04, 1994	2	1	Corrected Table reference
	2	2	Added indication of local use for Table 6
	1	34,34.1	Added new level specs, 115 & 116
Mar. 10, 1994	A	3	Moved Convective Precip to FOS
Mar. 24, 1994	1	4	Added gen. proc. for UVI (2), ditto for ozone analysis (two of them: 49,52)

[watch out replacing pages - pagination altered - following can refer to previously changed pages]

	1	27	Added parameter 24 for Ozone
	1	31	Added parameter 206 for UVI
Jun. 10, 1994	1	15	Added new 1.25 deg map No. 45
Jul. 01, 1994	1	3	Added NCAR to Table 0
	1	6	Added grid No. 45
	1	20	Added Arakawa grids for 29 km eta
	1	23ff	Noted location of pole on Lambert grids
	1	32	New parameters: 144-149
	1	39	New special level: 204
Jan. 02, 1996	all	all	A number of changes made in sections.
			Full revision/reprinting.
<u>March 10, 1998</u>	<u>0</u>	<u>all</u>	<u>A number of changes made in sections</u>
	<u>1</u>	<u>all</u>	<u>O and 1 and Appendix A (the later appended</u>
	<u>5</u>	<u>all</u>	<u>to Section 5. Full reprinting.</u>

Note: The most up-to-date version of this publication is available via **INTERNET** on the NOAA information center server.

Do an "anonymous FTP" into **nic.fb4.noaa.gov**, with a password of your e-mail address; check the /Readme file in the root directory for further directions.

GRIB Edition 1

INTRODUCTION

The World Meteorological Organization (WMO) Commission for Basic Systems (CBS) Extraordinary Meeting Number VIII (1985) approved a general purpose, bit-oriented data exchange format, designated FM 92-VIII Ext. GRIB (GRIdded Binary). It is an efficient vehicle for transmitting large volumes of gridded data to automated centers over high speed telecommunication lines using modern protocols. By packing information into the GRIB code, messages (or records - the terms are synonymous in this context) can be made more compact than character oriented bulletins, which will produce faster computer-to-computer transmissions. GRIB can equally well serve as a data storage format, generating the same efficiencies relative to information storage and retrieval devices.

Changes and extensions to GRIB were approved at the regular meeting of the WMO/CBS in February, 1988; additional changes were introduced at the CBS/WGDM/Sub-Group on Data Representation (SGDR) meetings in May 1989 and in October 1990. The 1990 changes were of such structural magnitude as to require a new Edition of GRIB, Edition 1, which this document describes. Further augmentations and interpretations were made by the SGDR in September 1993, with approval by the WGDM in February 1994. These changes did not result in a new Edition to GRIB, but did change some of the Tables, resulting a new Table Version number for them (2). The changes from Version 1 were mainly additions of new parameters or more precise definition of existing ones. Additional changes to the GRIB Tables were adopted at an Expert Meeting of the SGDR in September 1995, which now brings us to Table Version 3.

It is not anticipated that there will be any large-scale structural changes to GRIB for at least four to five years, or more. The SGDR is undertaking a thorough review of the present and future requirements that GRIB is supposed to satisfy. The plan is to design a major revision of GRIB capable of accommodating these requirements and more, without "straining" the structure of the data representation form. Some things are getting a little strained even now. See below.

Note: the Edition number is placed in the same location, relative to the start of the GRIB message, for all Editions. Thus, decoding programs can detect which Edition was used to construct a particular GRIB message and behave accordingly. This is useful for archives of messages encoded in earlier Editions or during transition periods. Of course, this requires that data centers retain copies of older editions of the code, and older versions of the parameter tables.

Each GRIB record intended for either transmission or storage contains a single parameter with values located at an array of grid points, or represented as a set of spectral coefficients, for a single level (or layer), encoded as a continuous bit stream. Logical divisions of the record are designated as "sections", each of which provides control information and/or data. A GRIB record consists of six sections, two of which are optional:

- (0) Indicator Section
- (1) Product Definition Section (PDS)
- (2) Grid Description Section (GDS) - optional

(3) Bit Map Section (BMS) - optional

(4) Binary Data Section (BDS)

(5) '7777' (ASCII Characters)

Although the Grid Description Section is indicated as optional, it is highly desirable that it be included in all messages. That way there will be no question about just what is the "correct" geographical grid for a particular field.

Most centers require bulletin headers to enable them to receive, identify, and switch messages; NCEP is no exception. The standard WMO abbreviated heading for GRIB is described in Appendix A.

In this documentation, certain symbols are used to clarify the contents of octets (groups of eight consecutive binary bits, "bytes" in American usage). If unadorned letters are used, they are symbolic and their meanings are described in the text; a decimal number is simply printed as is; a character or string of characters is represented inside single quote marks. International Alphabet No. 5, which is identical in its essential elements to the U.S. National Standard 7-bit ASCII, is used for character representation in the GRIB code.

Octets are numbered consecutively from the start of each section; bits within an octet are also numbered from left (the most significant bit) to right (the least significant bit). Thus an octet with bit 8 set to the value 1 would have the integer value 1; bit 7 set to one would have a value of 2, etc.

The numbering of Tables in the following text corresponds to the description of GRIB in the WMO Manual on Codes¹. Some additional tables not found in the WMO Manual are indicated by letters. These, generally, contain information unique to the NWS or NOAA.

A caveat: The Official International Documentation for GRIB is the just referenced Manual on Codes. This document is, in part, intended to be a guide to the use of GRIB and may not include all of the features currently found in the Manual. It does, however, represent the full set of features used by the National Weather Service, in particular in the AWIPS project, and by the National Centers for Environmental Prediction. The features described here are intended to be a completely consistent sub-set of the full WMO documentation; if there are any discrepancies the Manual on Codes is the final authority.

¹ World Meteorological Organization publication No. 306, Manual on Codes, Vol. 1, Part B, Secretariat of the WMO, Geneva, Switzerland, 1988, plus Supplements No. 1, 2, & 3 (with more to come)

This formula is the standard IBM representation for a single precision (real) floating point number. (Consideration is being given to using the IEEE floating point representation in the future, in a later Edition of GRIB.)

If second order (or "complex") packing is used (see the description of that later on) the internal value, X, will be made up of two values, a "local minimum value", X_i , and a "second order packed value", X_j . There will be one X_j for each grid point and a variable number of X_i values. This will all come clear later on when we get to the description of second-order packing.

What follows is a description, slightly simplified, of the process that one would go through to pack a (meteorological) field into a GRIB message, using "simple packing". It includes some explanations of why certain steps are taken, some of the consequences, and what choices have to be made. Some of the choices are interrelated; the relationships should be clear when the explanation is done. The additional features of complex or "second order" packing will be dealt with in a later section.

Given that a full field is available, the first step, if necessary, is to convert the units of the parameter into those shown in Table 2, the SI standard units, also known as the mks system. Some of the units may seem a little peculiar (kg/m^2 , for example, for precipitation - 1 kg/m^2 is equivalent to a water depth of 1 mm); others may seem inappropriate (Pa for pressure, for example, implies substantially greater precision than is typical in meteorological usage; inverse seconds are not nearly precise enough for divergence and vorticity) but they are all self consistent. The precision of the parameters, as actually packed in a message, can be set to any desired degree through the appropriate use of the power-of-10 ("D") scaling and the power-of-2 ("E") scaling. Just how this comes about will be described momentarily.

At this point there is a choice to be made. If it is desired to use a pre-selected bit word length for the packed variables, then just proceed on to the next step. However, if a variable bit word length is to be used, where the word length is adjusted to accommodate the data values, then it is necessary to undertake the power-of-ten scaling. The D value should be selected such that, when the original data, in the SI units of Table 2, is multiplied by 10^D , the integer part of the result will have enough precision to contain all the appropriate information of the variable. Anticipating things a little bit, the (scaled) value will be rounded to an integer as a part of the packing process; thus the "significant part" of the value of the variable has to be moved to the left of the decimal point prior to the rounding. Temperature might be scaled with $D=1$, thus changing the units to deci-degrees; pressure, on the other hand, might be scaled with $D=-2$, thus actually reducing the precision to hectoPascals (mb), a more reasonable meteorological precision; vorticity would be scaled up by $D=8$, and so on.

The second step in the packing operation is to scan through the field, which may or may not have been "D-scaled" at this point, find the minimum value of the parameter, and subtract that minimum - the reference value, R - from all the data points, leaving a residual of non-negative numbers. This step has two benefits. The first of these is convenience - making all the data points non-negative bypasses problems with different computer hardware that represent negatives in various ways: 1's complement, 2's complement, signed positive integers, whatever. The GRIB message is rendered just that much more machine independent by being non-negative throughout.

The second benefit is more consequential: it can result in a substantial compression of the bulletin size without any loss of information content. If a field has an appreciable bias away from zero, the residuals formed by the minimum removal operation will all be much smaller numbers than otherwise. Thus they will need fewer bits to contain them when they are, eventually, packed as integers.

The third step is simply to scan through the field of residuals and find the maximum value.

At this point another choice must be made, similar to the one made previously. This time, if a variable bit word length is to be used, then it is necessary to calculate how many bits (per word or per data gridpoint) are going to be needed to contain that largest data value, when the latter has been rounded to an integer. Recall that at the previous decision point, the variables were power-of-ten ("D") scaled such that a rounding operation will preserve all the significant part of the information. Discovering how many bits are needed is a simple scan through a table of powers of two, of course. The power-of-two-scaling is not employed and E is set equal to 0. Then go on to the fourth step.

If, alternatively, it is desired to use a pre-selected bit word length for the packed variables, the data must now be scaled, this time by a power of two (the "E" scaling), sufficient to either reduce the maximum value down to just fit into the available number of bits, or enlarge the value to just fit. This latter step takes care of the problem of small numbers where the precision is all in the fractional part of the number. How much precision is retained, for the eventual rounding, is a function of the preselected bit word length and the "typical" range, or maximum value with the minimum removed, of the particular variable. The choice of bit word length, which is made ahead of time, must be made with full knowledge of the characteristics of the particular variable that is to be packed and a prior assumption of how much precision needs to be retained for the largest likely value.

The fourth step is then to round all the values to integers, now that they have all been scaled to appropriate units, and pack them in the specified bit length words.

The last step is then to set up the various identification fields and put the GRIB bulletin in proper form. We shall turn to this "proper form" in the next section.

We have ended up with two alternate ways to construct a GRIB messages: a fixed bit word length method and a variable bit word length method. What are the relative advantages or disadvantages, or at least the differences, of one with respect to the other?

Message length: the fixed word length bulletins are always the same length, for a given parameter; the variable word length bulletins are, naturally, variable. The variation is driven by the range of the value of the parameter over the field (or the maximum value) which can change from day to day. Whether variations in message length is a problem or not depends on the computer systems used to work with the GRIB records.

Precision: The variable word length bulletins have a fixed and unchanging precision, determined by the "D" scaling. This assures that the same information content is available day after day. It is straightforward to change the precision in a familiar manner, that is, simply by orders of

magnitude, just by altering the D value. This comes at a cost, of course; increasing the precision by a power of 10 adds about 3.3 bits (average) to each data point in the message.

The fixed word length bulletins show a variable precision which is case by case data driven and is determined by the "E" (power-of-two) scaling that was used to fit the numbers into the available space. This can happen even with the same data, on the same date, but at adjacent grid areas. If one area shows a low variability and the neighboring one a high variability such that a different power-of-two scaling is needed in the two areas, then, unfortunately, the values on a common boundary will not be exactly equal after they are unpacked. This can be disconcerting and a cause for confusion. It will not happen if D-scaling (only) is employed. On the other hand, the variable precision can be viewed as a strength: a data field with a low variability will be encoded at a higher precision, thus preserving the character of the field; a high variability field will be represented with less precision, but that is not a problem as the small, and possibly lost, variations will not matter in the presence of the large ones. The precision of the encoded field can be increased by adding bits to the fixed word length, but the degree of change (a power of 2 for each bit) may not be as easy to deal with (or explain to people) as the simple order of magnitude change afforded by the "D" scaling method.

No matter which packing method was employed, a proper GRIB decoding program, that took account of the transmitted values of both "D" and "E", would return the correct unpacked numbers, regardless of which packing method was employed. It would be transparent to the user except for the questions of precision outlined above.

GRIB CODE FORM.

With the exception of the first four octets of the Indicator Section, and the End Section, all octets contain binary values. All sections contain an even number of octets; the variable length sections are padded with zero values as necessary. These extra bits must be accounted for in finding one's way through the sections; their content should be ignored.

SECTION 0: THE INDICATOR SECTION (IS)

The indicator section serves to: identify the start of the record in a human readable form, indicate the total length of the message, and indicate the Edition number of GRIB used to construct or encode the message. The section is always eight octets long.

Octet no.	IS Content
1-4	'GRIB' (Coded CCITT-ITA No. 5) (ASCII);
5-7	Total length, in octets, of GRIB message (including Sections 0 & 5);
8	Edition number - currently 1