12 IIPS - FA COMPRESSION OF WEATHER FORECAST DATA

Dietmar Saupe[1] *and Hannes Hartenstein*

Universität Freiburg
Freiburg, Germany

Werner Wergen

Deutscher Wetterdienst (DWD)
Offenbach, Germany

1. INTRODUCTION

This paper discusses the design and implementation of data compression applied to output from numerical weather prediction models. Such techniques will become essential when increased computational resources lead to much increased data volumes at weather forecast centers. The purpose of the data compression is to reduce the data storage requirements as well as transmission times and costs when serving remote outstations and customers with direct model output. Important issues to be considered are precision of the decoded data, the encoding and decoding speed, the complexity of the coding techniques, and the compatibility with existing standards such as GRIB.

As more powerful computers become available the precision of local weather forecasts can be improved by using higher resolution in space and time for weather simulations. For example, at the German Weather Service (DWD), typically grids of 200 by 200 points are used for each of 20 model levels in the vertical. It is anticipated to replace the current mainframe computer in order to allow to increase the resolution to 800 by 800 by 50. This will blow up the resulting data file sizes by a factor of 40. Even though storage devices are becoming less expensive there is obviously a need for compression of these large amounts of data.

We have investigated the prospects of compression techniques which should satisfy the following criteria:

- There should be only one method applied for all different types of data such as precipitation, pressure, wind, etc.
- The compression should be lossy to achieve a considerable compression rate, but there should be a guaranteed accuracy of the encoded data, which must be under user control.
- The compression achieved should not be surpassed by a significant amount when using more sophisticated methods.
- The compression should be easy to program allowing maintenance also by non-experts of data compression.

2. OVERVIEW

We have considered transform coding using the discrete cosine transform (the JPEG standard), wavelet based methods and predictive coding. The last of these methods appears to have the most advantages in the context of our application. It proceeds in three steps briefly described below.

1. **Quantization.** The various types of forecast fields have different ranges and different distributions within these ranges. For example, the range of height data might increase from 800 m at 1000 hPa to 3000 m at 50 hPa. Furthermore, the fields are only known within the limits of the observation errors. These errors can serve as guidelines for precision requirements imposed on the compression algorithms. For height data, the maximal observation errors are assumed to be 5 m at 1000 hPa and 25 m at 50 hPa. Therefore, the data can be uniformly quantized with quantization errors within these tolerances. For other fields, e.g., precipitation, precision requirements vary with precipitation amounts. The compression should preserve dry areas and not produce "packing drizzle." For large amounts, the requirements are less severe. Thus, under such circumstances a *logarithmic transformation* can be applied before uniform quantization. After the uniform or logarithmic quantization with user specified tolerances the

[1]Corresponding author address: Dietmar Saupe, Universität Freiburg, Institut für Informatik, Am Flughafen 17, 79110 Freiburg, Germany, email: saupe@informatik.uni-freiburg.de. This paper appears in Proc. 12th Intern. Conf. on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology, American Meteorological Society, Atlanta, January 1996.

data can be losslessly encoded allowing exact reconstruction of the quantized values.

2. **Optimal prediction.** The suitably quantized data already require less storage. However, there still is a lot of spatial coherence that can be exploited for further data reduction. This means that data items can be estimated well from surrounding data values. It suffices to store the difference between the predicted values and the true ones. The differences that need to be stored are much smaller than the original values and therefore need fewer bits to store. This scheme is known as *lossless predictive coding.* The prediction at a given data point is given by a linear combination of a few reconstructed data points that have already been coded in the process. The coefficients can be chosen optimally in order to minimize the sum of the squared prediction errors.

3. **Entropy coding.** The storage requirements can further be reduced by standard entropy coding such as Huffman or *arithmetic coding* (we chose arithmetic coding).

All of these steps (except for the logarithmic quantization) are common practice in data compression. The emphasis in this paper is on the application in the meteorological context. The gold standard in meteorology is GRIB (see FM 92 GRIB) but it differs with respect to all of the above three aspects. The quantization is always uniform and can be specified only indirectly by allocating a fixed number of bits per data value (16 bits are used at the DWD). The actual precision within the GRIB file then is given implicitly by the number of bits specified and the range of the data. Neither spatial coherence nor coding redundancy are considered in GRIB. While this seems to have been the appropriate solution for the coding and file format of weather data at the time of the GRIB design (1985), it is becoming clear now that enhancements such as those described here are useful today and will be indispensable in the future.

3. RESULTS

We have experimented with five types of data:

1. Height data from the "Global-Modell" with a resolution of 1.5 degrees. This is one of the most widely used fields. At DWD, it is stored for 12 pressure levels in the vertical and, thus, occupies considerable storage space in the data banks.

2. Wind fields from the 14 km "Deutschland-Modell (DM)." As this is a vector field, special attention is required. Also, at low wind speeds, the compression should not revert the wind direction.

3. Precipitation from the DM. The packing should not introduce spurious rainfall ("packing drizzle").

4. Diffusion coefficient (DM). This parameter varies over several orders of magnitude.

5. Roughness length (DM). There are considerable differences between land and sea, which should be preserved by the compression algorithm.

The optimal prediction and the entropy encoding steps are standard and identical in all cases. Details on these methods can be found, e.g., in Rabbani and Jones (1991) and Witten, Neal, and Cleary (1987). In order to portray the adaptations necessary for the compression of the various types of data files mentioned above we discuss the first three of them.[1]

3.1. Pressure

For the encoding of isobaric surfaces at pressure levels ranging from 50 hPa to 1000 hPa we adapt the required precision at the pressure level. The low-pressure surfaces are very high, have larger tolerances for the physical measurements, are relatively smooth as compared to the high-pressure surfaces and, thus, can be compressed more efficiently. Table 1 shows the results.

In order to compare the GRIB format with the results given in Table 1, we calculate the number of bits needed to represent a single data value at a predetermined accuracy in the GRIB format. This number can be directly compared to the bps (bits per sample) column in Table 1. For example, the range of the 400 hPa chart is 1465 m and the maximum error introduced by data compaction should not exceed 12 m. Using GRIB representation with binary scaling, we have to pack the values, so that the quantization levels are 16 m wide, corresponding to an error bound of 8 m. With this setting a total of $\lceil 1465/16 \rceil = 92 \leq 128$ quantization levels are necessary resulting in an optimal bit rate of 7 bps for GRIB. In contrast, with our method we only need 1.1 bps, corresponding to a compression ratio of 6.4 over the optimized GRIB as above. In practice, non-optimized GRIB files are used with a bit rate of 16 bps leading to an actual compression ratio of 14.7.

[1] A full version of the details of the encoding and the experiments can be found in Saupe and Hartenstein (1995).

isobaric surface hPa	range m	accuracy m	file size kB	bit rate bps	comp. ratio
1000	753	5	5.3	1.5	4.7
950	765	6	4.9	1.4	5.0
850	814	6	4.6	1.3	5.4
700	980	7	4.2	1.2	5.8
500	1261	10	3.9	1.1	6.4
400	1465	12	3.8	1.1	6.4
300	1741	13	4.0	1.1	6.4
250	1932	14	3.9	1.1	6.4
200	2128	16	3.4	1.0	7.0
100	2485	22	2.1	0.6	11.7
50	2875	25	1.7	0.5	14.0

Table 1: Results of compression of various data sets containing different constant pressure surfaces. Each data set consists of 121×240 samples. The file size, given in kilobytes (kB), represents the size of the compressed data set; bps stands for the average number of bits used to represent a sample value (the bit rate). The compression ratio is computed by comparing the bit rate of the fifth column with the bit rate of the optimal GRIB representation (see the text for details).

3.2. Precipitation

With the precipitation data it is not adequate to apply a uniform quantization as for the case of pressure data because at low precipitation values a much higher resolution is desirable than for large precipitation values. We propose to quantize that data using a threshold $S > 0$. Precipitation less than the threshold will be regarded as zero and coded by a special symbol. For precipitation values x greater or equal to the threshold S we use a logarithmic quantization, i.e.,

$$x \mapsto \text{round}\left(\frac{10\ \log_{10}(x)}{R}\right)$$

where $R > 0$ is an appropriate constant, a logarithmic resolution, measured in dB (decibel). This constant determines the precision of the encoding as follows.

At the decoder the number $10\ \log_{10}(x)/R+\delta$ with a rounding error of $|\delta| \leq 0.5$ will be received. Thus, the reconstructed value $\tilde{x}$ of x will be

$$\tilde{x} = 10^{\log_{10}(x)+\frac{\delta R}{10}} = 10^{\frac{\delta R}{10}}\, x.$$

Then the relative error is

$$\left|\frac{\tilde{x}-x}{x}\right| \leq 10^{\frac{R}{20}} - 1 \approx 0.115\, R.$$

For example, with a logarithmic resolution $R = 0.1$ dB we obtain a relative error bound of 1.15%. In our experiments precipitation is measured in kg/m^2 and the threshold was set to $S = 0.001$ kg/m^2 (see Table 2).

data type	range kg/m^2	accuracy	file size kB	bps
large scale	0 – 111	11.50%	4.4	3.0
precipitation		1.15%	9.2	6.4
convective	0 – 230	11.50%	6.3	3.7
precipitation		1.15%	11.4	7.8

Table 2: Compression results for precipitation data sets from the “Deutschland-Modell.” Each data set contains 109×109 values. Because of the logarithmic quantization no absolute accuracy can be given. Instead a relative error bound in percent can be guaranteed.

data type	range	accuracy	file size	bps
wind angle	36^o	5.0^o	3.1 kB	2.1
		0.5^o	7.4 kB	5.1
wind speed	13 m/s	0.50 m/s	1.9 kB	1.3
		0.05 m/s	5.5 kB	3.8

Table 3: Compression results for a wind data set from the “Deutschland-Modell.” Again, the data set consists of 109×109 samples.

Here we cannot compare the performance directly with GRIB, because GRIB does not support logarithmically scaled data files.

3.3. Wind

Wind data is usually given as a two-dimensional array of 2D vectors (two files for wind velocities in West and North directions) while the precision requirements are more intuitively based on wind direction and wind speed. For example, one might require the wind direction to be precise within ± 5 degrees, and the wind speed within ± 0.5 m/s. Therefore, we propose to first transform the original data into polar coordinate form. Then each component can be quantized uniformly with its prescribed precision. Note that the predictor for the angular data can be designed to take advantage of the 2π-periodicity of the space of angles. Table 3 summarizes the results for one representative data set from the “Deutschland-Modell.”

3.4. Other data

We complete our experiments with studies on data sets of the “Deutschland-Modell” for the roughness-length and the diffusion coefficient, see Tables 4 and 5.

data type	range	accuracy	file size	bps
diffusion	0.0003 –	$5.5 \cdot 10^{-5}$	7.7 kB	5.3
coefficient	0.067	$5.5 \cdot 10^{-6}$	14.4 kB	9.9

Table 4: Compression results for a diffusion coefficient data set from the "Deutschland-Modell."

data type	range	accuracy	file size	bps
roughness	0.000001 –	11.5%	3.8 kB	2.6
length	6.8	1.15%	8.7 kB	6.0

Table 5: Compression results for a roughness length data set from the "Deutschland-Modell."

4. CONCLUSION

The results from the last section show that a considerable amount of storage can be saved by compression, even at conservative requirements regarding accuracy. These savings will even increase as the spatial resolution of the data goes up because this will cause the prediction errors to decrease dramatically which reduces the number of symbols to be encoded and makes the entropy coding more efficient.

We have tested our routines also on data stemming from extreme weather events. The performance of the compression method is very similar to the cases listed above.

The encoding and decoding times are small. For example, on one of our UNIX workstations (SGI Indy with R4000 processor) these times were typically at most 0.5 seconds including the file I/O.

We have compared the performance of our method with standard JPEG transformation coding and with a wavelet based coder. Neither one gives significantly better results. Moreover, these methods are more complex, the encoding times are longer, and guarantees for given data accuracy are not immediate.

Two other desirable properties of the method proposed in this paper are: When a data set is encoded, then decoded and then encoded again, the two encoded files are identical, i.e., no additional loss or change is made to the data. Secondly, if two data sets that differ at some pixels are encoded, then their decoded versions differ at exactly the same locations. These properties cannot easily be shared by other approaches such as JPEG or wavelets.

5. SUMMARY

We have presented a simple, yet efficient method for weather data compaction. It is based on quantization followed by optimal linear prediction and entropy coding of the prediction errors. The user specified accuracy can be guaranteed either absolutely or in terms of relative change using logarithmic quantization. Future work will include an implementation and testing of this method for a general compressor of GRIB type files.

6. REFERENCES

FM 92 GRIB World Meteorological Organization – No. 306 Manual on Codes, Vol. 1, Intern. Codes, Part B – Binary Codes.

Rabbani, M., Jones, P. W. *Digital Image Compression Techniques*, SPIE Optical Engineering Press, Bellingham, 1991.

Saupe, D., Hartenstein, H. *DWD-Projekt: Kompression von Wettervorhersage-Daten*, internal report, Institut für Informatik, Universität Freiburg.

Witten, I. H., Neal, R., Cleary, J. G, *Arithmetic coding for data compression*, Communications of the ACM 30,7 (1987) 520–541.