

C. Wunsch · D. E. Gunn

A densely sampled core and climate variable aliasing

Received: 5 August 2002 / Accepted: 19 May 2003 / Published online: 10 July 2003
© Springer-Verlag 2003

Abstract Undersampled records are susceptible to aliasing, in which a high frequency appears incorrectly as a lower one. We study the sampling requirements in a core taken from Rockall Trough using bulk density, P-wave velocity, and magnetic susceptibility as measured on an automated system. At 2-cm spacing (approximately 33 years in this core), all variables show a characteristic red-noise behavior, but with a spectral slope that is sufficiently weak so that significant aliasing appears to be present. P-wave velocity shows the largest potential corruption, probably owing to the weaker spatial averaging present in the sensor. Approximately 50% of the apparent low-frequency energy is aliased in all variables at some frequencies in both quiet and active regions of the record. In this core, a sampling interval of 0.2 cm appears to be “safe” in both active and quiet portions of the core, aliasing little or no energy, except in the P-wave record. For cores of different duration, sampling interval, and measurement type, the considerations will be identical, the importance of the problem depending directly upon the shape of the overall spectrum describing the entire frequency (wavenumber) range of physical variability.

Introduction

Studies of climate change are based largely upon sampled depth/time series of various physical variables.

C. Wunsch (✉)
Department of Earth, Atmospheric and Planetary Sciences,
Massachusetts Institute of Technology, Cambridge,
MA 02139, USA
E-mail: cwunsch@mit.edu
Fax: +1-617-2534464

D. E. Gunn
Challenger Division for Seafloor Processes,
Southampton Oceanography Centre,
Empress Dock, Southampton, SO1 3ZH, UK

“Samples” range from near-point measurements in a core to local averages within the core. Averages are obtained in different ways: by summing many point measurements, by melting an ice core over some finite interval or by a transducer with a finite footprint in the core. Many measurements made in cores are quite onerous (e.g., Foraminifera) and so the core is sampled at comparatively infrequent intervals. Other measurements can be done automatically (e.g., magnetic susceptibility) and in principle, the core could be sampled arbitrarily densely, up to the limits of sensor dimension.

With any record that has been sampled (that is, reduced from a hypothetical continuous variable to a finite set of discrete numbers), a major issue is always whether the sampling rate is adequate. The issue of adequacy is dealt with in all textbooks on time-series analysis, but perhaps a very brief recapitulation is helpful here. Consider a record consisting of a single pure sinusoid, with circular wavenumber (or frequency) $\kappa_1 = 1/L_1$

$$f(z) = \sin\left(\frac{2\pi z}{L_1}\right) \quad (1)$$

where z is a continuous variable. Assuming a uniform sedimentation rate, only a scale factor distinguishes between depth z and time t . Then, it is easy to confirm that if it is sampled, uniformly, at intervals $\Delta z < L_1/2$, the wavenumber visible in the samples is still κ_1 . To the contrary, if $\Delta z > L_1/2$, the periodicity of the samples is changed to an apparent wavenumber of

$$\kappa_a = \left| \frac{1}{L_1} - \frac{n}{\Delta z} \right| \leq \frac{1}{2\Delta z} \quad (2)$$

where n is an integer such that the inequality is satisfied. This effect is most familiar in the stroboscope, where a shutter is opened at just the right time to render a moving, periodic object, like a wagon wheel, apparently fixed, by choosing $\kappa_a = 0$. The physical statement is that a sinusoid that is not sampled frequently enough will still appear to be a sinusoid, but at a lower frequency. That is, it masquerades as, or “aliases” to an incorrect

frequency. The highest estimable wavenumber is at the Nyquist frequency, $\kappa_N = 1/(2\Delta z)$; any energy present at higher wavenumbers will appear artificially at some lower wavenumber given by Eq. (2).

Aliasing may be less familiar in the context of so-called continuum stochastic processes, in which one can refer to energy in a particular band of frequencies, but without the process being periodic. Such energy has precisely the same sampling behavior as in Eq. (2), appearing as a lower-frequency, non-periodic energy. (Textbooks, e.g., Priestley 1982, should be consulted for the general Fourier representation of stochastic processes.) The consequences of failing to sample records sufficiently often to leave the dominant energy at the appropriate place in frequency space are not always completely appreciated.

A general tool for the description of climate and other systems is the sample power density spectrum, $\tilde{\Phi}(\kappa)$, in which one estimates the frequency content of the record from the measurements. The resulting structure of the estimated spectrum is both a useful description, and a powerful diagnostic of the physical processes going on. Use of the spectrum depends upon it being an accurate estimate; undersampled records, i.e., those having Δz (or the corresponding Δt) too large, may greatly distort the apparent form. As in any statistical estimation problem, it is helpful to distinguish between the theoretical function, $\Phi(\kappa)$, and the values which are actually estimated, $\tilde{\Phi}(\kappa)$. The two are not identical, but only the latter is available here.

In situations for which sampling is comparatively easy (e.g., a physical oceanographer attempting to measure currents with a moored instrument), the stakes are sufficiently high, that one seeks reassurance about the measurement strategy. Such reassurance would normally be obtained by a preliminary experiment in which, for some fraction of the record, “oversampling” was done at a rate $\Delta z'' \ll \Delta z$, such that one could confirm that there was little or no energy at wavenumbers exceeding $1/2\Delta z$; otherwise the sampling strategy would be modified so as to be adequate.

A small number of studies exists on the aliasing question in the paleoclimate literature, including Pisias and Mix (1988) who called attention to its importance, and Wunsch (2000). The latter was directed at explaining

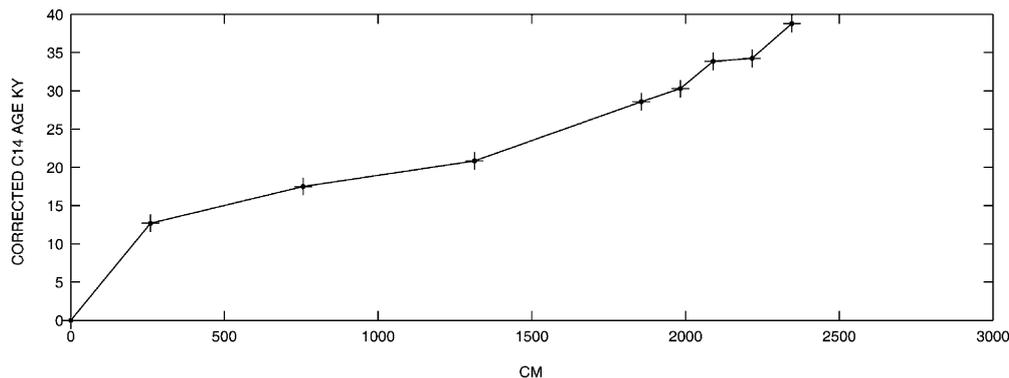
a mystifying, sharp peak in a Greenland ice core which appears to be a simple alias of the seasonal cycle. Here, we ask a more general question.

Materials and methods

Marion Defresne core MD95-2006 was taken in the Rockall Trough; its stratigraphy was discussed by Knutz et al. (2001). This core was selected somewhat arbitrarily for analysis for two reasons: it was readily available to us at Southampton Oceanography Centre (SOC), and had been moderately densely dated by radiocarbon methods. The extent to which specific conclusions drawn about sampling this core can be generalized is unknown; each record must be studied separately until widely applicable rules emerge. We offer the present results only as an illustrative example. A depth–age relationship for the core is depicted in Fig. 1. As is common, an initial record at 2-cm (about 33-year intervals on average) sampling was used (Fig. 2).

Three variables were measured: P-wave (acoustic) velocity, magnetic susceptibility, and bulk density. As with many properties recorded in cores, exactly what these variables represent in the climate system remains obscure. For present purposes, they are generic representatives of the sampling problems raised by any variable thought to reflect climate change. Because the three physical variables are measured with three different techniques using an automated multi-sensor logging system (Gunn and Best 1998), we must briefly describe the sensors. Magnetic susceptibility is measured with a transducer (Dearing 1999) which produces an average over an elliptical area with a minor diameter of 0.5 cm along the axis. P-wave velocity is measured over a circular area of diameter 2 cm. Bulk density is measured using a 2-mm diameter gamma-ray collimator; spreading between the collimator and the sample increases the sensor footprint to 5 mm at the point of measurement. The extent to which the measurements are uniformly averaged over these sensor footprints can be determined either from knowledge of the sensor details or from the data themselves. With a uniform 1-cm average, for example, displacing the sensor by 0.1 cm would generally produce only a slight change in measured value—because

Fig. 1 Carbon-14 dates (from Knutz et al. 2001) for core MD95-2006 as a function of depth. For purposes of this paper, a linear interpolation between dated points (+) was used, and no data more recent than 14,000 a B.P. were included



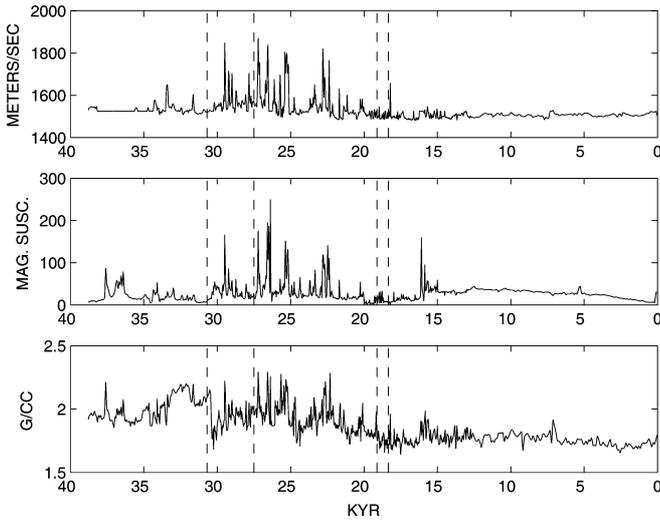


Fig. 2 Core MD95-2006 data sampled at 2-cm (averaging about a 33-year) intervals. Dates are from Fig. 1. The *upper panel* shows measured P-wave velocity; the *middle panel* is magnetic susceptibility (SI units) $\times 10^{-5}$ (that is, measured values are $O(10^{-5})$). The *lower panel* is bulk density. The relatively active and quiet regions between 27.5 and 30.7 ka B.P. and between 19.11 and 18.35 ka B.P. were resampled at 0.1-cm intervals

of the large overlap. However, a 0.1-cm average could vary greatly 0.1 cm away. To the extent that the measurement is weighted toward the center of the sensor, one would expect to see larger changes between neighbors. Sensor “transfer responses” from point disturbances could be determined from experiment; to our knowledge, no such experiments have been done.

The degree to which a transducer reduces high wavenumbers present in the record involves the shape and size of the transducer relative to the physical variations. Suppose that z is the depth variable in the core, and that y is the corresponding horizontal dimension across the core. If a physical variable, η , present in the core is dominantly in vertical wavenumber $k = 2\pi\kappa$, then

$$\eta(z, y) = A \cos(2\pi\kappa z) \quad (3)$$

independent of y . Suppose the transducer is circular, of radius r_1 . Then, the measurement is

$$\tilde{\eta}(z_i) = \int_0^{r_1} \int_0^{2\pi} w(r, \theta) A \cos(2\pi\kappa r \cos \theta) d\theta dr \quad (4)$$

Here r, θ are polar coordinates centered on the core at position z_i midway in the cross section, and $w(r, \theta)$ is a weighting function describing the transducer response (Fig. 3). One can expect that it does not, to a first approximation, depend upon θ . We have no information about their behavior as a function of r for any of the sensors used here and so, as a guideline, will simply assume that $w = 1/(2\pi r_1)$, that is, a uniform average. Equation (4) then transforms to

$$\tilde{\eta}(z_i) = \frac{A}{r_1} \int_0^{r_1} J_0(2\pi\kappa r) dr \quad (5)$$

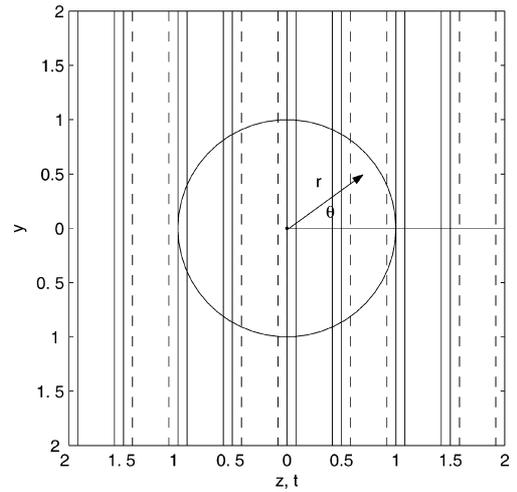


Fig. 3 Averaging done by a cylindrical sensor on a core whose axis is z , for depth, or t for time, with laminae running across the core in the y direction. *Dashed contours* are -0.5 , and *solid contours* are $0.0, 0.5$

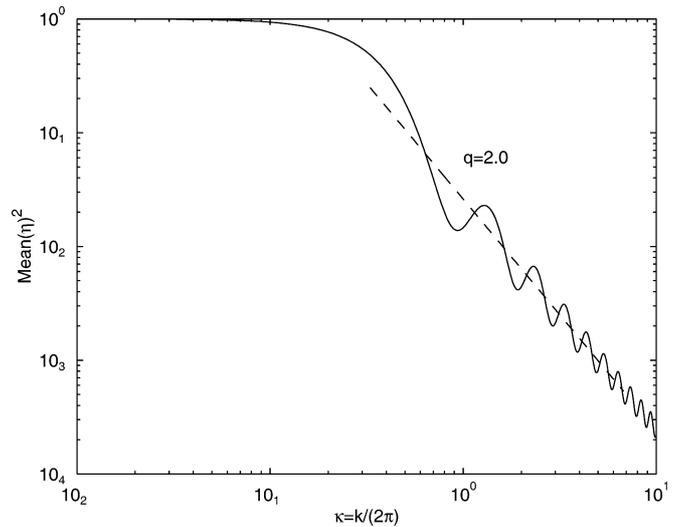


Fig. 4 Attenuation of the energy in a sinusoid as a function of wavenumber when averaged over a circle of unit radius. Plotted as variance (square of the value of η)

where J_0 is the Bessel function. This integral can be evaluated in terms of Struve functions, but is more easily found numerically (Fig. 4) and shows a -2 power-law drop-off for wavenumbers k beyond about $k=1$, in a physically plausible result. Thus, if the measurement is equivalent to a uniform integral over the circular transducer, and if the transducer radius is 1 cm, one expects to see the spectrum drop at least as fast as κ^{-2} beyond $k = 2\pi\kappa \approx 1/\text{cm}$. This argument assumes, as is commonly true, that the spectrum of the continuous record does not increase with increasing wavenumber. A similar analysis can be done for an elliptical sensor such as the one used for magnetic susceptibility, but the result is qualitatively the same. The P-wave velocity measurement is somewhat problematical because the sensor

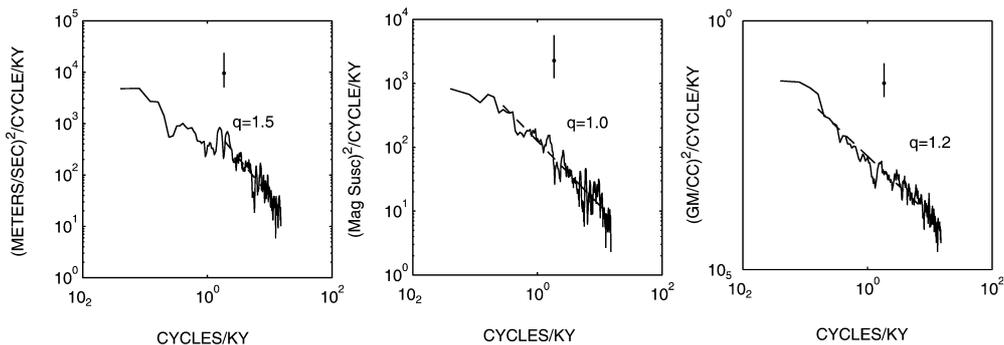


Fig. 5 Dimensional power density spectra of core MD95-2006 before the Holocene (before 14 ka B.P.) from 2-cm (approximately 33-year intervals) sampling. The oldest data used were at 38,794 a B.P. The *left panel* shows P-wave velocity, the *middle panel* the magnetic susceptibility (SI units $\times 10^{-5}$), and the *right panel* the bulk density. There is only a faint suggestion of a steepening of the spectra near the high-frequency cutoff, and the slopes near $q=1$, suggesting significant aliasing. Power-law fits were done over frequency bands (*dashed line* best-fitting straight line) selected visually to be representative of the high-frequency behavior. Here and elsewhere, an approximate 95% confidence limit is shown

records only the first arrival, which may well emanate from only a fraction of the area occupied by the sensor. Thus, the degree of spatial averaging in this measurement will vary with depth, and no simple analysis is available in this case. In the results below, we will see some apparent consequences of this lack of averaging.

Results

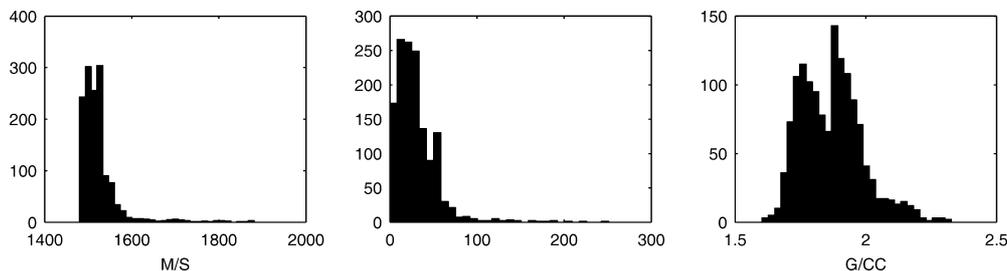
Spectra of the full record

The computed spectra of the three variables as sampled at 2-cm intervals are shown in Fig. 5. Data more recent than 14 ka were not used in these spectral estimates, as the Holocene portion of the core was probably disturbed (E.J.W. Jones, personal communication 2001). A least-squares power-law fit of the form $\Phi(s) = As^{-q}$ is displayed for each of the variables over the frequency interval spanned by the dashed line. All records have high-frequency power laws close to s^{-1} , with the spectra approaching white noise at the lowest estimable frequencies. As reassurance that one has adequately sampled any record, one hopes to see a steeper drop-off in energy with s as the highest frequency is approached, or the achievement of a white-noise level there. The latter

would indicate that the high-frequency energy is at the round-off noise or “least-count” level—where the measurement represents only the random rounding error of the discretization (e.g., Bendat and Piersol 1986; Bomar 1998). In the present case, neither behavior is seen. A rule of thumb is that if $q \geq 2$, one can sub-sample with little error but, for smaller values of q , there will be significant aliasing. It is possible that the record contains no significant energy at wavenumbers above one cycle/4 cm, the highest resolved wavenumber at 2-cm sampling, but one should remain skeptical until the high wavenumber (high frequency) energy is actually determined. That none of the power density spectral estimates shows the $q=2$ or steeper power-law behavior expected from a true spatial average suggests either that the record is being undersampled, or that high-frequency energy increases sufficiently fast so that the spatial averaging is inadequate to suppress it. Notice that none of the spectra shows any sign of the ice-core peak at $s = 1/1,470$ year = 0.68/kiloyear and ascribed by Wunsch (2000) to a seasonal alias. In a deep-sea core, the absence of such a peak is consistent with some combination of bioturbation, an approximate 1-cm spatial averaging corresponding to about 16 years, and the general suppression of annual cycle signals in the non-equatorial deep sea.

In general, the variables measured here are non-Gaussian; empirical frequency functions are shown in Fig. 6. Many natural processes are nearly Gaussian, most commonly interpreted as having resulted from the summation of large numbers of uncorrelated events. For such processes, it is easy to show that the statistics of the data are completely described by the record mean and the spectral densities. Statistical properties of the sample means, spectral estimates, and other useful properties, such as the incidence of extreme events, are fully worked

Fig. 6 Empirical frequency functions from the core at 2-cm sampling. *Left to right* P-wave velocity, magnetic susceptibility, and bulk density. All are qualitatively non-Gaussian (see text)



out and well understood. The present records display heavy tails, meaning that there is more to the records than the information found in the means and spectra (one requires so-called higher-order, or polyspectra, for completeness), and the probability of rare events in the record is very different from what a Gaussian process would display. The extent to which the records are also simultaneously non-stationary is a very difficult question to answer without a-priori knowledge of the true probability density functions, and we will not further pursue the problem here.

A highly-sampled energetic interval

The spectral results above correspond to the average behavior over the entire record. Visually at least, the data in Fig. 2 exhibit a strong non-stationarity, with quiescent intervals contrasting with much more active ones. Although such behavior is not a rigorous demonstration of a true statistical non-stationarity (on long time scales, climate may have a bimodal probability density, one which was temporally invariant), it is still useful to ask whether the sampling requirements on the records differ in the active and quiescent periods. This does not mean that the quiescent intervals are less important, and they are discussed separately below.

An energetic region between 1.34 and 1.5 m in the core was resampled at 0.1-cm intervals (an approximate 1.4-year sampling interval), shown in Fig. 7. For this short core segment, the age–depth relationship is a linear one, and the frequency domain analysis could equally well have been in the depth coordinate rather than in time. A conversion to time was made simply to render the result more physically interesting. Because the

spatial position of the samples is accurately known, there should be no “jitter” effect in the sampling for results referred to a depth coordinate reference frame (see Wunsch 2000, and the references therein).

In the lower panel of each pair in Fig. 7, the estimated time rate of change of the records, using the two different sampling intervals, is illustrated. Depending upon exactly where the 2-cm samples are placed, one can hit or entirely miss a major event in the record. Whether such brief events are an accident of the geological record, such as an unusual but not significant ice-rafting event, or whether they imply a real, short-lived climate event reflecting the long tails of the frequency functions is one of the imponderables of the subject (most events do not occur in all three variables simultaneously).

Time rates of change are shown because almost all theoretical constructs for physical variables, such as temperature or ice volume, involve not the variable but its time derivative. Thus, most theories are schematically of the form

$$\frac{d\theta}{dt} + \dots = F \quad (6)$$

so that if θ is temperature, F would be a heating/cooling rate. Alternatively, if it is ice volume, F would include things such as precipitation, ablation and the like (Wunsch 2003). To make use of such theories, one must be able to estimate the *rates of change* of the variables. Sparsely sampled records can miss much of the structure of the time derivatives.

The differences between the records at the two sampling rates can also be examined in the frequency domain. Figure 8 displays the normalized spectral density estimates for the three variables, as well as of their rate of change. Note that in the remainder of this paper,

Fig. 7 The core as sampled at 0.1-cm intervals (approximately 1.43 years) in a visually active region (*solid line*). The *first* (*uppermost*) *panel* is the P-wave velocity, and the coarsely sampled version (2 cm) of the same variable, displaced arbitrarily in the vertical to distinguish them visually (*dash-dot line*). The *second panel* is the time derivative of the two *upper curves*, calculated from a simple finite difference. Time differencing exaggerates the high frequencies, but these would be important in any theory of the variable. The *third and fourth panels* are the same as the two *uppermost panels*, except for magnetic susceptibility. The two *lowermost panels* are for the bulk density

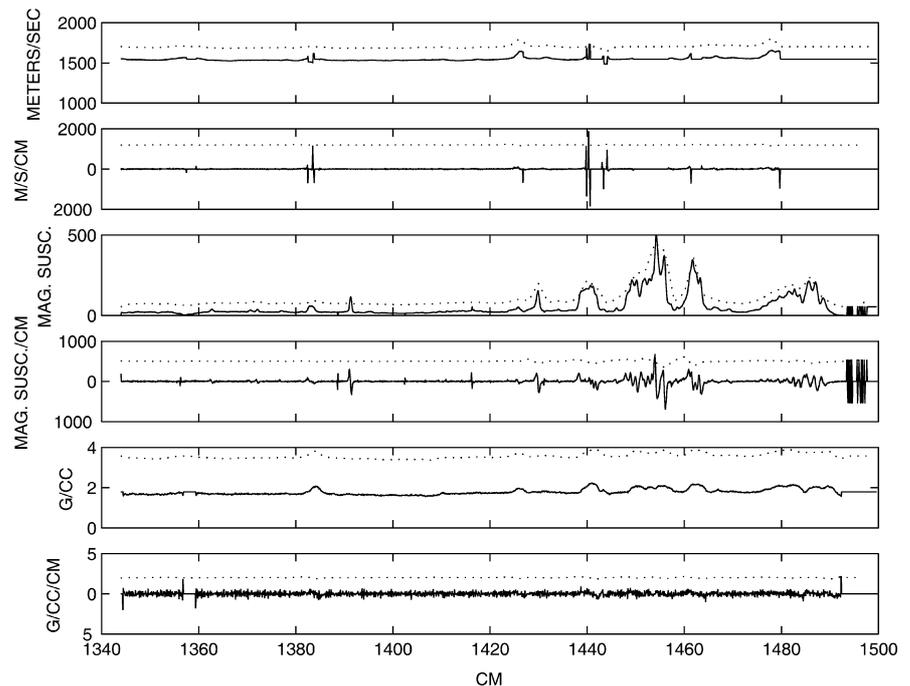
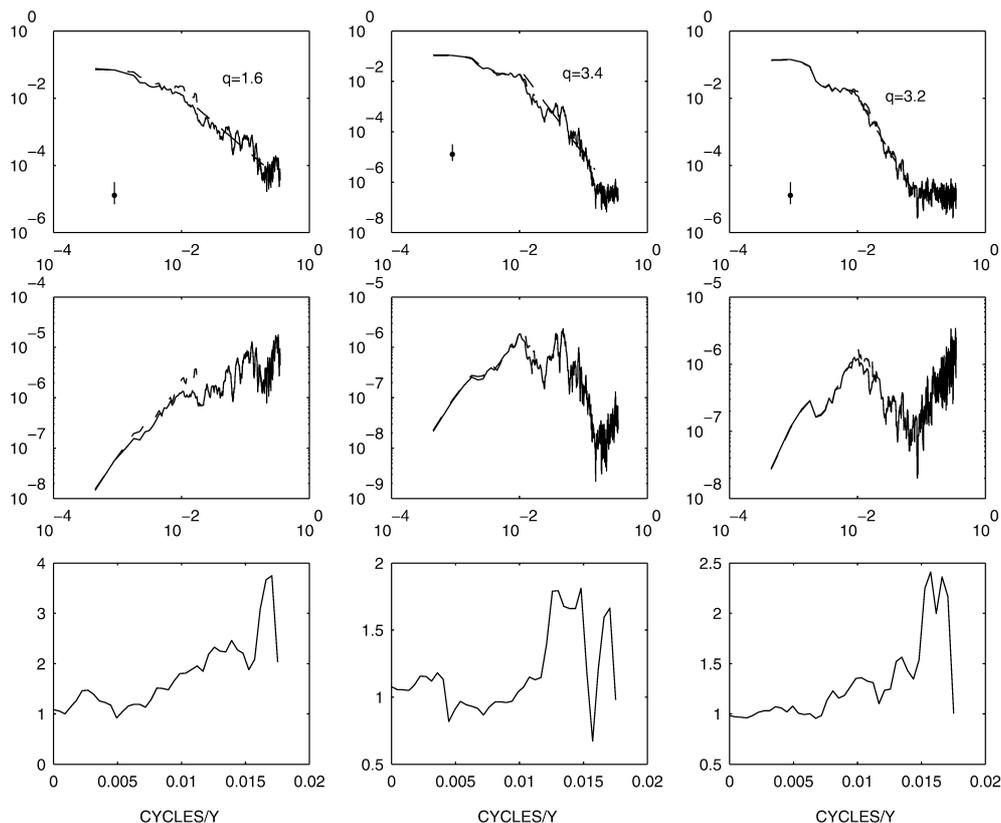


Fig. 8 Normalized non-dimensional spectra of P-velocity, magnetic susceptibility, and bulk density (*left to right*) from the active region of the core at 0.1-cm sampling (approximately 1.43 years). The *upper panel* is the spectral density estimate, the *middle panel* is the spectral density estimate of the time rate of change of the variable. In the *lower panel* is shown the ratio of the energy densities in the coarsely sampled record, to that for the finely sampled one. Note that the frequency scale for the spectral ratio is linear, whereas it is logarithmic in the other plots, and all are in cycles/year. Values in the ratio plots significantly greater than 1 are indicative of an aliasing problem. That the problem increases with frequency is consistent with their being a red-noise structure in the true spectrum at low frequencies. *Dotted lines* are the aliased estimated spectral densities



time is measured in years, rather than kiloyears, and the spectra are displayed non-dimensionally so that the integral under the curves is unity. This form is used because the shape of the spectrum is of primary concern, and not the absolute level. (The variance of the densely sampled and sub-sampled records remains the same).

As noted above, in an adequately sampled record, the spectrum ideally becomes flat (white) at a low level at the very highest accessible frequencies. Conventionally (Bendat and Piersol 1986, p. 339), the interpretation then is that no actual signal is present at these frequencies, and one has reached the irreducible spectral floor given by the round-off (quantization) level of the digital signal. Magnetic susceptibility appears to have reached this level near 0.15 cycles/year and bulk density near 0.08 cycles/year. Both bulk density and magnetic susceptibility show power laws steeper than $q=3$ at high frequencies, and so some further degree of sub-sampling could be tolerated. The part of the spectrum corresponding to the quantization noise level produces a spurious “blue” power spectrum in the rate of change, which should be suppressed in any quantitative use.

It is less clear whether the P-wave velocity sampling has reached the quantization level; this result is perhaps not a surprise, given the comments above about the lack of spatial averaging in that measurement. The velocity measurement also exhibits a flatter power law at high frequencies, and thus is probably significantly aliased.

A highly sampled quiet interval

The records of a quiet core interval are shown in Fig. 9, and the corresponding power-density estimates are in Fig. 10. Both magnetic susceptibility and bulk density show the high-frequency, least-count white-noise floor near 0.3 cycles/year. By contrast, P-wave velocity again has a behavior indistinguishable from a power law up to the highest estimated frequency. These behaviors generally mimic the spectra in the active region, although the high-frequency behavior is quantitatively somewhat different in the two cases. Note that the apparent low-frequency power laws differ quantitatively in the aliased and non-aliased cases, an issue if spectra are used to test theoretical ideas about climate change.

Discussion

The overall conclusions to be drawn here are simple. The three parameters, which are sufficiently easy to measure so as to permit modest oversampling, suggest that a 2-cm spacing (corresponding to an approximate 33-year interval, on average, over this core) produces significant low-frequency aliasing. By “significant” is meant that 50% or more of the apparent low-frequency energy is spurious. This aliasing means that the apparent low frequencies are more energetic than they should be. Such contamination has many consequences. Consider, for example, the problem of determining whether two

Fig. 9 In the *upper panel* the P-wave velocity from the finely sampled core (*solid line*) as well as the coarsely sampled record (displaced vertically, *dash-dot line*) are shown. The *middle panel* shows the record for magnetic susceptibility, and the *lower panel* for bulk density. The 0.1-cm sampling interval corresponds to a time-sampling interval of about 0.6 years. The record is taken from between 19.1 and 18.4 ka B.P. in Fig. 2

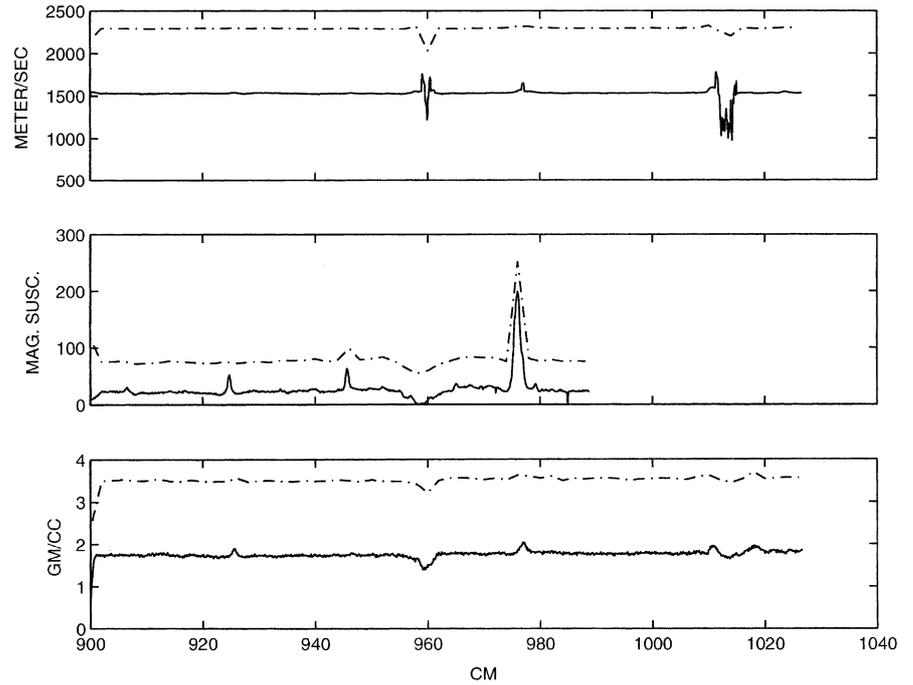
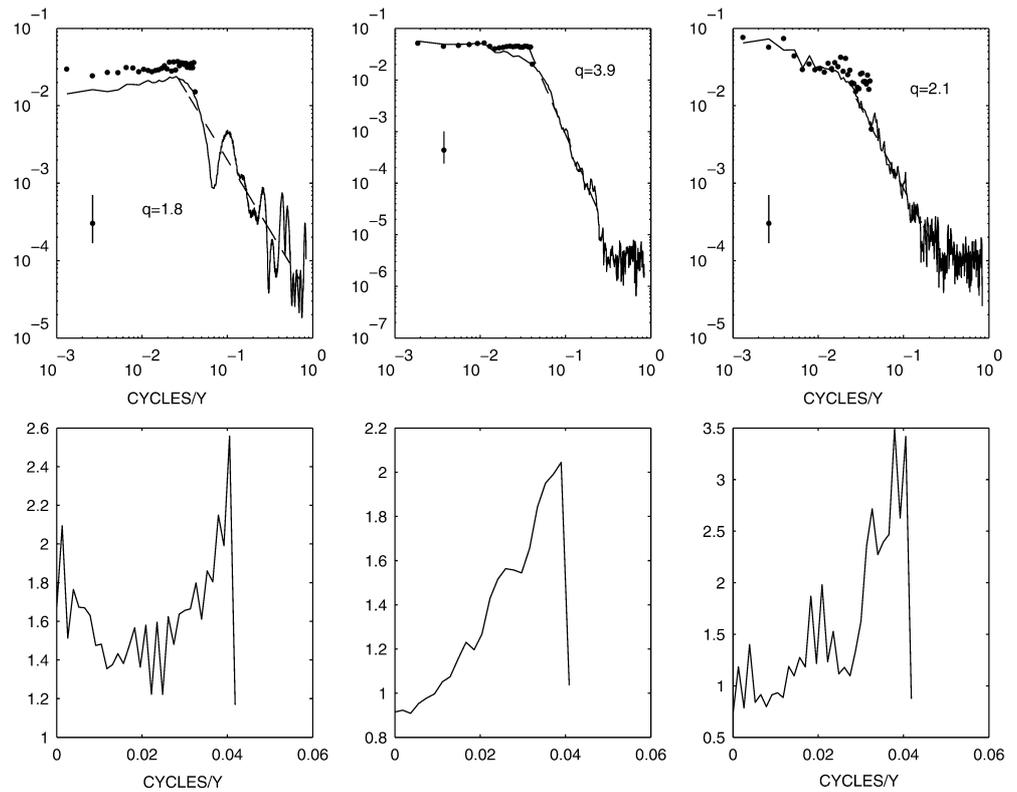


Fig. 10 Power density spectra of P-velocity, magnetic susceptibility, and bulk density (*left to right*) from the quiet interval of the core and sampled at 0.1-cm (approximately 0.6 years) intervals. The *lower panels* are the ratio of the energy in the coarsely sampled spectrum to that in the finely sampled one, in the region of overlap



records are coherent with each other at low frequencies. Suppose, for simplicity, that the true coherence is $\gamma = 1$ (perfect coherence), but that the aliased high-frequency energy is completely incoherent (not necessarily true), and that only one record is susceptible to aliasing. Then, since the coherent power computed is only 50% of the apparent power in the aliased low frequencies, the

apparent coherence would be reduced to $\tilde{\gamma}^2 = 0.5$ or $\tilde{\gamma} = 0.7$. Should both records contain 50% aliased energy, then $\tilde{\gamma}^2 = 0.25$. If, as would be expected, the true coherence is less than 1, then it can be wholly swamped by the aliased energy.

For bulk density and magnetic susceptibility in this core, with sensors like those used here, it would

appear adequate to sample at a rate such that $1/(2\Delta t) > \approx 0.2$ cycles/year, i.e., with a corresponding interval of $\Delta t < 2.5$ years. Or, in the spatial domain, so that $\Delta z < 0.2$ cm in both active and quiet regions. In the quiet parts of the core, the steeper spectra would permit a coarser resolution with tolerable aliasing. The extent to which the conclusions drawn here are applicable to other cores remains unknown, and must be explored for each record and variable separately. Sampling with such high density is prohibitive for many variables, especially those which are not automated. Depending upon the underlying spectrum, one can use a variety of strategies to minimize the problem, including so-called burst sampling, in which closely spaced observations are made in “bursts” at comparatively rare intervals. Such a strategy is useful only if there exists a gap in the spectral energy intermediate between the low frequencies of interest, and the high frequencies which must be suppressed (not true for MD95-2006). At a minimum, if an adequate sampling procedure is unavailable, one must estimate the degree to which energy has been aliased, as in some cases it may well be the dominant noise element.

We have used the terminology here of “low” and “high” frequencies, but these are of course relative terms. For a core representing 100×10^6 years of time, energies estimated at low frequencies are markedly different than in a core spanning only 30,000 years. In the former core, samples may be spaced apart at intervals Δt not of years but of thousands of years or longer. What matters is the spectral shape in whatever time intervals are appropriate, and indeed the entire discussion here is best done in completely non-dimensional time or depth units. The result, however, is somewhat abstract, and we leave to the reader the (simple) translation of the terminology of low and high frequencies to whatever is best in a particular situation.

None of the variables used here is close to having a Gaussian probability density. Without a physical understanding of the actual probability density, it is difficult to determine whether the underlying processes are non-stationary, or whether the apparent times of statistically variable behavior are just the result of heavy-tailed probability densities.

Acknowledgements This work was begun while C. Wunsch was a visitor at the Southampton Oceanography Centre and at University College, London. Thanks are owed to E.J.W. Jones for stimulating this study. P. Huybers made some helpful comments.

References

- Bendat JS, Piersol AG (1986) Random data: analysis and measurement procedures, revised edn. Wiley-Interscience, New York
- Bomar BW (1998) Finite wordlength effects. In: The digital signal processing handbook. CRC Press, Boca Raton, pp 3-1-3-18
- Dearing J (1999) Environmental magnetic susceptibility. Using the Bartington MS2 system. Chi Publishing, Kenilworth, UK
- Gunn DE, Best AI (1998) A new automated nondestructive system for high resolution multi-sensor core logging of open sediment cores. *Geo-Mar Lett* 18:70–77
- Knutz PC, Austin WEN, Jones EJW (2001) Millennial-scale depositional cycles related to British Ice Sheet variability and North Atlantic paleocirculation since 45 kyr B.P., Barra Fan, U.K. Margin. *Paleoceanography* 16:53–64
- Pisias N, Mix A (1988) Aliasing of the geologic record and the search for long-period Milankovitch cycles. *Paleoceanography* 3:613–619
- Priestley MB (1982) Spectral analysis and time series. Vol 1, univariate series. Vol 2, multivariate series, prediction and control. Academic Press, London
- Wunsch C (2000) On sharp spectral lines in the climate record and the millennial peak. *Paleoceanography* 15:417–424
- Wunsch C (2003) The spectral description of climate change including the 100 ka energy. *Climate Dyn* DOI 10.1007/s00382-002-0279-z