# Building a Statistical Model to Predict Reactor Temperatures

Carl Scarrott and Granville Tunnicliffe Wilson*

October 3, 2000

**Abstract**

This paper describes the various stages in building a statistical model to predict temperatures in the core of a reactor, and compares the benefits of this model with those of a physical model. We give a brief background to this study and the applications of the model to rapid on-line monitoring and safe operation of the reactor. We describe the methods, of correlation and two dimensional spectral analysis, which we use to identify the effects which are incorporated in a spatial regression model for the measured temperatures. These effects are related to the age of the reactor fuel and the spatial geometry of the reactor. A remaining component of the temperature variation is a slowly varying temperature surface modeled by smooth functions with constrained coefficients. We assess the accuracy of the model for interpolating temperatures throughout the reactor, when measurements are available only at a reduced set of spatial locations, as is the case in most reactors. Further possible improvements to the model are discussed.

*Keywords: SPATIAL PREDICTION; TWO DIMENSIONAL SPECTRUM; LINEAR MIXED MODEL*

## 1 Background and objectives

The Magnox nuclear reactors were built with a designed lifetime of about 30 years, which they have safely out-lasted. For optimal efficiency the reactors are operated close to temperature limits that are determined by safety considerations, and these limits have been gradually lowered over the years. A Magnox reactor contains thousands of fuel channels and the limits are applied to the channel gas outlets temperatures (CGOTs) where the coolant gas is at its highest temperature. In most reactors, however, only a fixed subset, typically 10%, of the CGOTs, are monitored. These measurements provide a good indication of the temperature cross-section of the reactor. The reactor is operated so that these are as high as possible, subject to the operating limit.

*Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, UK.

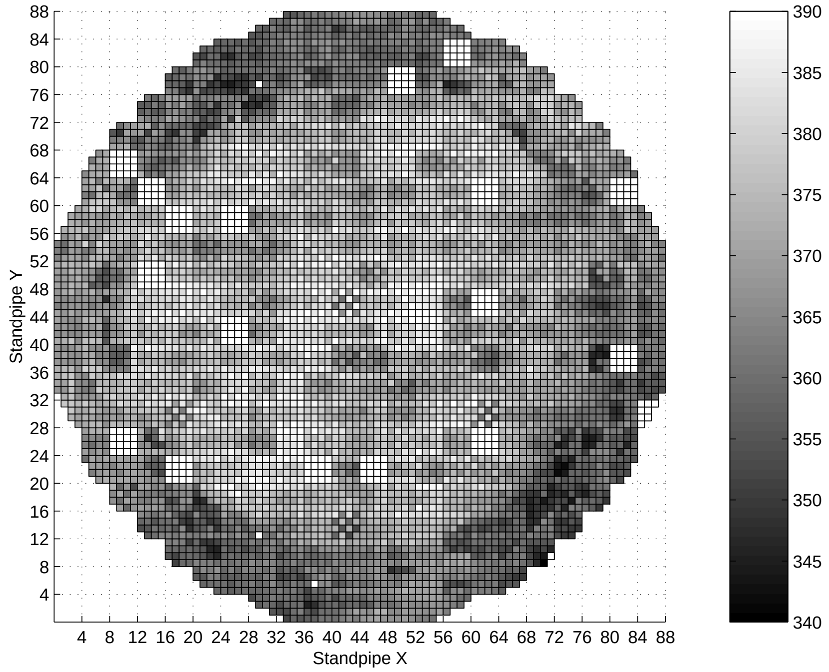


Figure 1: Scan CGOT measurements in °C for Wylfa Reactor 1, March 1997. White squares are missing values.

There are two reactors at Wylfa where, exceptionally, all the CGOTs can be measured in a *scan*. Figure 1 is a representation of a scan of Wylfa reactor 1 (R1), measured in °C. We discuss this in more detail in section 3, but remark here that the figure reveals local areas which are generally hotter or cooler, though there is still substantial variation within them. Operator control can only raise or lower the general temperature level in large sectors of the reactor; individual CGOTs cannot be separately manipulated. Consequently the measured CGOTs are generally distributed well below their maximum, as shown for Wylfa by the histogram in Figure 2a. In reactors where only a subset of CGOTs are monitored, temperature limits are applied to these values. Figure 2b shows the distribution of measured CGOTs for Dungeness R1 and R2, where the CGOTs are monitored only for every third channel in both directions, a pattern we refer to as the $3 \times 3$ subgrid. The upper limit of 390°C is lower than for the Wylfa reactors because some of the unmeasured off-grid CGOTs are likely to lie above the imposed limit.

The prime objective of developing a model for CGOTs is therefore the prediction, or interpolation, of unobserved CGOTs from the sub-grid of monitored values. A statistical method for interpolating these CGOTs, using only the monitored values to estimate the local mean temperature surface, is presented by Logsdon and Tunnicliffe Wilson (2000), referred to hereafter as LTW. Our present objective is substantially to reduce the uncertainty in these predictions by incorporating the known effects of reactor geometry and fuel age into a spatial regression model.

The development of a statistical model for this purpose also has the value of pro-

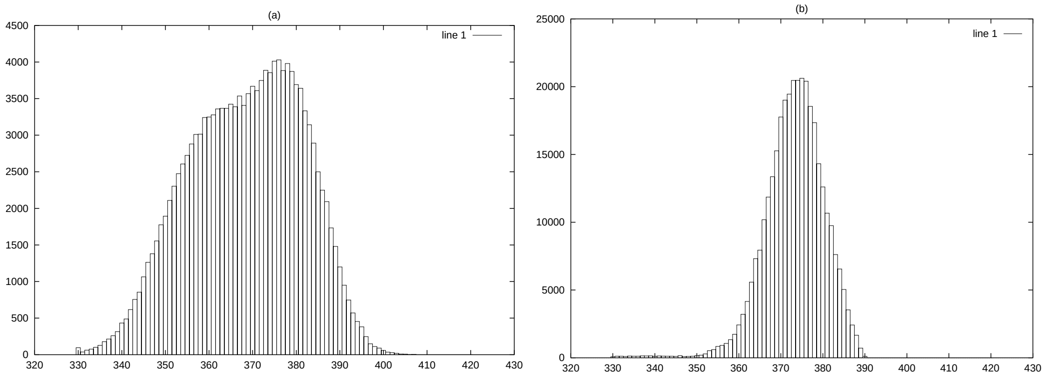


Figure 2: Histograms of measured CGOTs at Wylfa and Dungeness reactors, collected over a period of one year.

viding an independent approach for comparison with a physical model, developed and used by the nuclear industry, for predicting reactor conditions. This model uses the nuclear properties of the core, such as the fuel, moderator and control rod insertions. It has a wide range of uses such as planning of reactor refueling and reactor fault studies for safety assessment. The statistical model, because it is developed to fit a large set of reactor operating data, is much more empirical than the physical model. A comparison of the predictive accuracy of the two models can serve to improve each of them. For example, the statistical model has identified the importance of some effects of reactor geometry which have not yet been incorporated into the physical model. A simple *post-hoc* modification based on this effect can then substantially improve the physical model predictions.

A further benefit of the statistical model is its potential as an on-line tool for operator assessment of reactor scans. Visual display of the fitted temperature surface and residuals assists validation of the scan by revealing any local CGOT anomalies which may be ascribed, for example, to thermocouple defects. This can be done promptly following the scan because fitting the model requires only one or two minutes. Carried out routinely, this could make an important contribution to data quality and risk assessment.

# 2 Data and Main Effects

We shall describe the data used for modeling the Magnox reactor Wylfa R1. In this reactor all CGOTs may be measured in a scan taken over a short time period. We will build our model using all this data, then re-fit the model using a $3 \times 3$ sub-grid of CGOTs. By this means we can assess the accuracy of the statistical model for predicting, or interpolating, the remaining temperatures using only the sub-grid. At the other Magnox reactors, where data is only available for a sub-grid of channels, interpolation of the remaining temperatures is one of the main objectives. The exercise with Wylfa is intended to develop and validate the models we use for interpolation at the other reactors.

The Wylfa reactors (R1 and R2) each have 6156 fuel channels, and are physically the largest reactors in the world. Figure 1 is a representation of one temperature scan of R1. Each square represents a fuel channel, which form a regular grid over the reactor cross-section. The colour of the square represents the temperature over the range 340-390°C. The channels are numbered in Cartesian form; the distance between channels is about 9 inch, and is called a pitch unit. The reactor is kept very stable during the process of the scan, so the data can be considered a snapshot of the reactor temperatures at one particular time-point.

The central part of the scan, a circular region excluding channels within 10 pitch units of the periphery, corresponds to the *inner* or *flattened region*. The aim is, by operator control and the refueling plan, to maintain CGOTs close to the upper limit throughout this region. Outside this region, is the *outer* or *unflattened region*. Here the CGOTs are somewhat lower.

The whole reactor is structured according to a 4 by 4 grouping of channels, called standpipes, which are evident for example at $(41-44, 13-16)$, $(61-64, 29-32)$ and $(69-72, 29-32)$. The core can only be accessed by ports above each standpipe, so most reactor operations are carried out via these. In particular the fuel is replenished one standpipe at a time. The chequer-board effect visible within some standpipes is mainly caused by the fuel being replaced in a chequer-board pattern. The fuel in one part of the chequer-board is replenished at one time-point and some time later the complementary part is refueled. The fuel reactivity and power output is low when first inserted in the reactor, increasing quickly to a fairly constant level over its remaining lifetime, tailing off a little at the end. When the first half of the chequer-board is replenished the new fuel will be relatively cold compared to the old, causing the chequer-board pattern.

The standpipes coloured in white indicate missing measurements. Because the measurement process is also carried out by standpipe, if the equipment develops a fault, the data may be lost for the whole standpipe. The pattern of missing data varies between scans.

The main explanatory variable for the variation of CGOTS about their local mean level, is the fuel age, measured by the irradiation to which it has been exposed. In each fuel channel there are several fuel cans, which will have differing irradiation. Figure 3 is a representation of the average fuel irradiation in each of the channels, measured in giga-watt days per tonne (GWD/Te). The oldest fuel is in white and the youngest fuel is in black. The irradiation cannot be directly measured until the fuel is removed from the reactor at the end of its lifetime, so we use an estimate determined from past reactor operating conditions. The standpipe refueling patterns are clearly visible. For every reactor the estimated fuel irradiation age is known for *all* channels, so is an important variable for prediction of unmeasured CGOTs.

We shall see clearly that the CGOTs are directly affected by the age of fuel in the corresponding channel. They are, however, also known to be affected indirectly by the age of fuel in surrounding channels by the mechanism of neutron diffusion. Modeling of this effect is an essential feature of the physical reactor model.

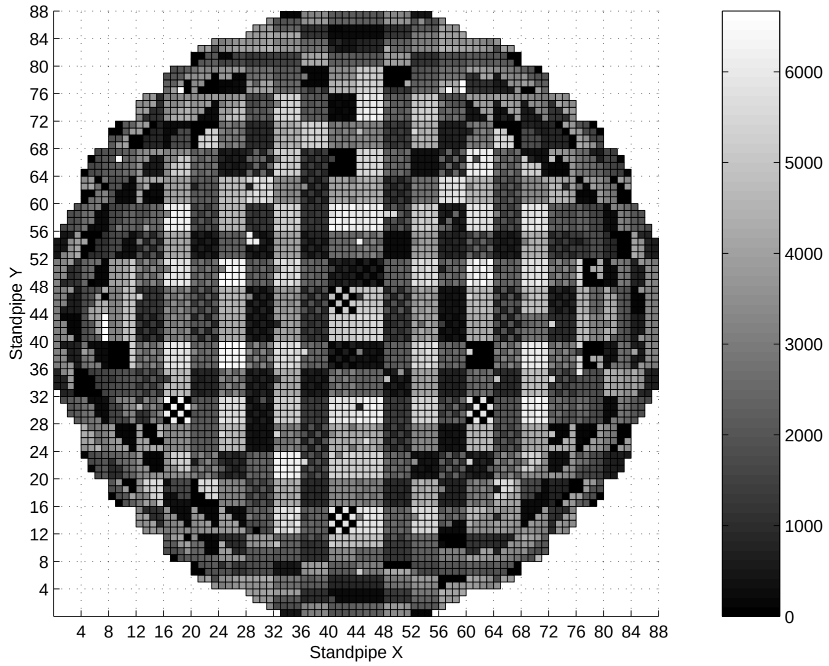


Figure 3: Channel Average Irradiation in giga-watt days per tonne (GWD/Te) for Wylfa Reactor 1, March 1997. Old fuel is in white, new in black. Note the standipe pattern of refueling which is reflected in Figure 1.

The variations in CGOTs may also be explained in part by effects ascribed to the reactor geometry. For any given reactor these will vary little with time and will mostly be regular in structure across the reactor, reflecting its basic design. The magnitudes of several of these effects were assessed using a simple flow model in an industry study of 1978, but the statistical verification was limited. Most of these effects are not included in the main physical reactor model. We remark here on the origin of a selection of these geometry effects. Figure 4 shows the channel brick layout within a standpipe. The fuel channels are vertical holes through a column of graphite bricks which are of two sizes, hexagonal and square in an alternating layout. The larger hexagonal bricks absorb more neutrons and hold heat better than the smaller square channels. This causes an approximate 2.5°C difference in the temperature distribution in an alternating pattern across the reactor cross-section.

Along the central two rows of channels the corners have been cut out of the bricks giving circular 'interstitial' holes, some, but not all of which contain control rods and fixed absorbers. These holes are blocked at the bottom, but coolant leaks into them, between the graphite bricks, mostly in the lower part of the reactor where the differential pressure is greatest and the coolant is at a lower temperature. The effect is to cool the adjacent CGOTs, but the larger hexagonal bricks have a greater surface area in the interstitial holes and are cooled more. The result is east to west banding of the temperatures in rows of width two pitch units, of magnitude about 4.5°C, and an additional alternating effect of approximately 1°C in the central two rows of standpipes. The control rod and fixed absorber holes are also of different diameters, which leads

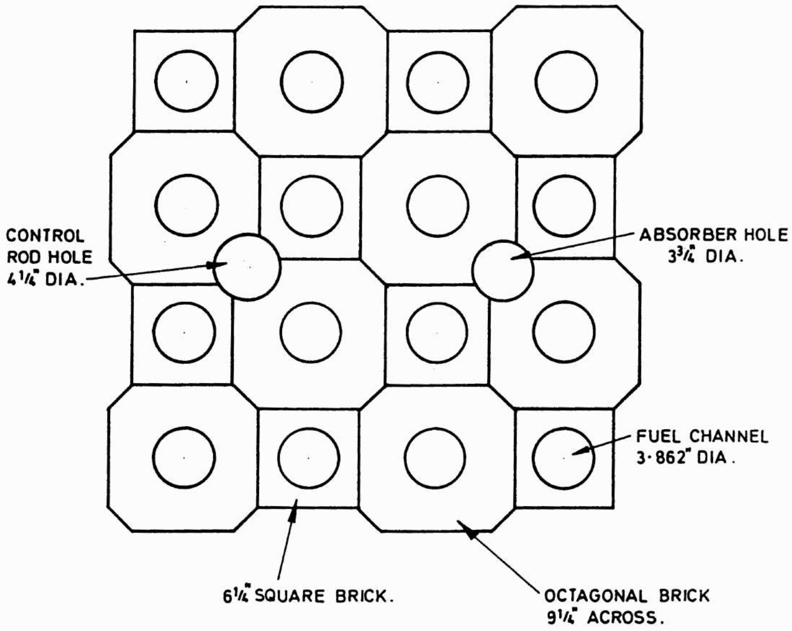


Figure 4: Geometry of a reactor standpipe.

to a further differential effect on fuel channels adjacent to such rods, of about 1°C. We have quoted the magnitudes of these effects from the report of the 1978 study. Using appropriately patterned regressors, our model was able to verify the presence of these and other geometric effects, such as control rod type and insertion, and estimate them more precisely.

The effects of fuel age and reactor geometry on the CGOTs generally fluctuate rapidly over a distance of a few pitch units. A large component of the CGOTs is, however, the relatively slow variation in their local mean level across the reactor. This is difficult to predict from reactor variables because of the sensitivity of the reactor to the insertion of control rods. A small movement of one rod could lead to a relatively large change in the general level of CGOTs over a substantial part of the reactor. In the physical model this sensitivity is managed in part by normalising the total predicted reactor power to its measured value, and the use of other similar, semi-empirical, *power-shaping* methods.

In our spatial regression model we shall include separate terms to model this slowly varying or low frequency component of CGOTs. This is purely empirical, but justified on two counts. First, it is necessary to ensure accurate estimation of the effects of fuel age and geometry. Secondly, when unknown CGOTs are predicted from values observed on a sub-grid, these terms are essential to estimate the correct local level. In an earlier analysis, which included no explanatory variables apart from one minor geometry effect, LTW estimated this low frequency component using the non-parametric technique of kernel smoothing, see for example Hastie and Tibshirani (1990). In our spatial regression model we shall use instead a large set of smooth regressors, capable of fitting any smooth surface. To avoid over-fitting we shall give them random coefficients with variances specified to ensure that any rapidly varying components have relatively small magnitude. Their overall scale will be controlled by a single parameter chosen to minimise the cross-validation sum of squares.

# 3 Statistical Exploration

We describe how the effects described in the previous section are investigated by correlation and spectral methods, starting with the fuel age variable. Figure 5 is a plot of

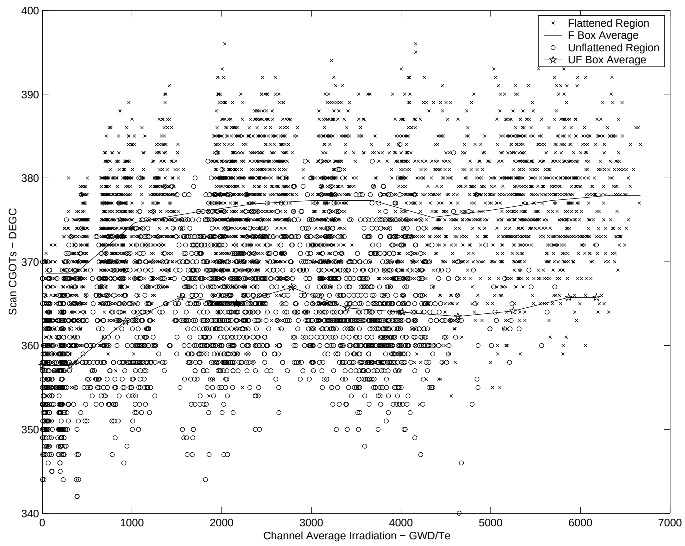


Figure 5: Channel temperatures in °C against channel average irradiation in GWD/Te with 10 box moving average to show general behaviour. Inner region channels are crosses and outer region are circles.

the channel temperature against channel average irradiation. Data from the inner and outer regions are indicated by different symbols. The box moving average shows the mean response, which is fairly flat apart from the rise at the start. The pattern in the two regions is similar apart from a difference in level.

The effect of fuel irradiation on CGOT seen in this figure is not strong, because it is masked by the other effects. For example, in any region having a concentration of fuel with lower irradiation, evidence of a lower level of CGOTs would tend to be offset by operator control attempting to restore the level.

In order better to study the relationship between the fuel irradiation and temperatures we can partly correct for such effects. We do this by looking at the contrast between each CGOT and the surrounding mean temperature level, and relating this to the similar contrast for fuel irradiation. We call this pre-whitening, following a similar procedure in time-series analysis, see Box and Jenkins (1976, p 379). The resulting contrasts have much lower spatial correlation, so appear more random. We use the kernel smoothing developed by LTW to estimate the local level and subtract this from each CGOT, applying exactly the same procedure to the fuel irradiation. Figure 6 is a plot of the pre-whitened channel average CGOTs and fuel irradiation. Much stronger evidence of a predictive relationship is clear, with a linear correlation of over 0.6 for both regions. The moving box means are similar for the two regions and show the expected increase of temperature with irradiation early in the life of the fuel.

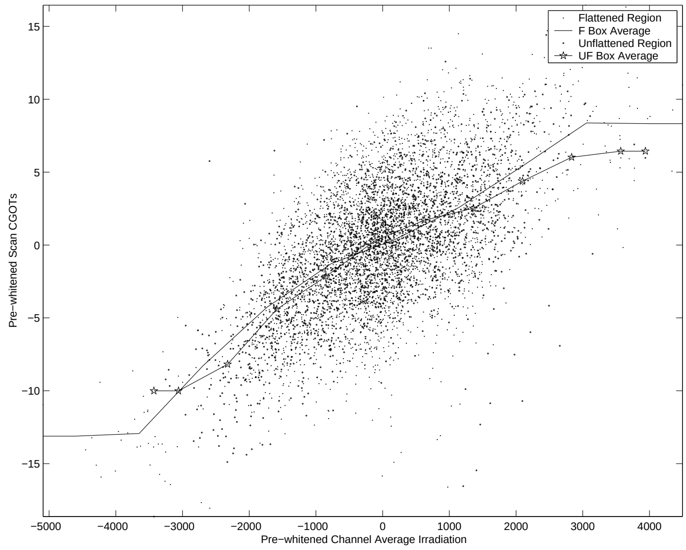


Figure 6: Pre-whitened channel temperatures against pre-whitened channel average of fuel irradiation with 10 box moving average to show general behaviour.

It is natural to progress to spatial spectral analysis. This decomposes the variance of the data into components associated with a range of spatial periods. We use the multi-taper methods illustrated for example in Zhang (1994) and originating with Slepian (1964). Details of our approach are given in Scarrott and Tunnicliffe Wilson (2000). We present here selected parts of the analysis to communicate its usefulness. Figure 7 shows the two dimensional power spectrum of the CGOTs. A two-dimensional taper was applied to the data but no spectrum smoothing was carried out. A similar spectrum formed for the fuel irradiation is shown in Figure 8. Most striking are the sharp peaks in Figure 7. For several of these there is no corresponding peak in the fuel irradiation. They are caused by the regular geometric effects of the reactor design. Their position indicates their period, typically 2 and 4 pitch units, and their orientation, whether east-west or north south or both. After removing these peaks, as described in the next section, lower peaks were seen which could be attributed to the pattern of fuel irradiation, because corresponding peaks were clearly visible in the spectrum of that variable.

To confirm the relationship between the fuel irradiation and CGOT scan, we carried out a cross-spectrum analysis between them, using multi-tapers to perform the spectral smoothing. We remark that the cross-coherency was generally high, but we present here only the estimate of the spatial impulse response function shown in Figure

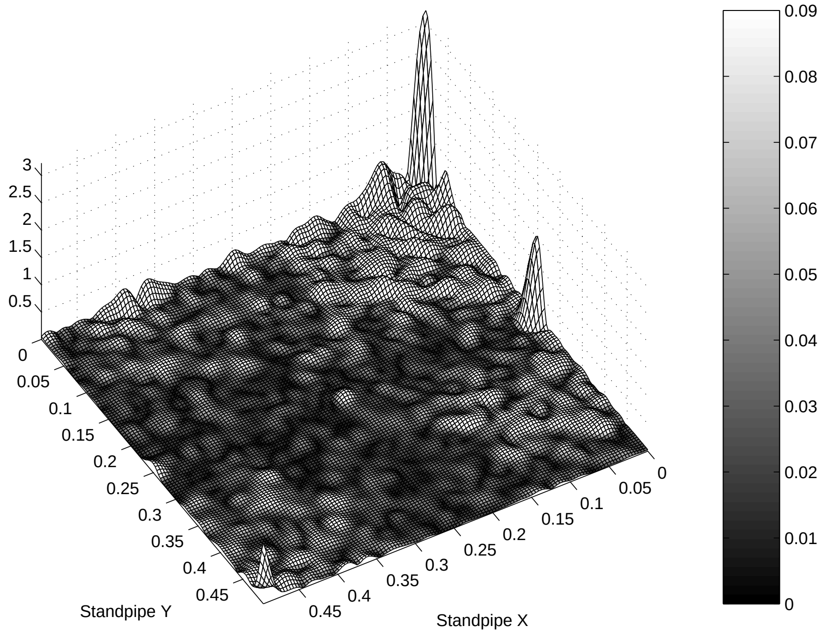


Figure 7: Power spectrum of CGOTs calculated using a single taper, and square root transformed to reveal detail.

9. This shows what would be the pattern of CGOTs arising from a single channel of fuel, with high irradiation age, surrounded by fuel of a uniformly low age. The central peak indicates the increase in the CGOT in the central channel. Adjacent to the peak are slightly negative values. More distant are small but positive values extending out for several pitch units. The more distant effect measures the increase in surrounding CGOTs due to the increased stimulation of fission by neutron diffusion from the central channel. The adjacent negative effect arises because the older fuel absorbs more neutrons. This reduces local reactivity, off-setting the effect of neutron diffusion. We have used a robust randomisation procedure to determine the significance of this response pattern. The central peak is highly significant, the surrounding effects only marginally so. Cumulatively however, the non-central effects can be important.

# 4 Spatial regression modeling

We first developed a regression model over a central square region of the reactor having sides of length 62 pitch units. The choice of such a region initially facilitated the use of standard two dimensional spectral analysis.

Taking as response the measured CGOT $y_{ij}$ at pitch coordinates $i$ and $j$, the model is

$$y_{ij} = f\left(x_{ij}\right) + g_{ij} + h_{ij} + e_{ij} \tag{1}$$

The first explanatory term, $f$, is a non-linear function of the irradiation age, $x_{ij}$. The second represents the combined effect of regressors, each with its own linear coefficient,

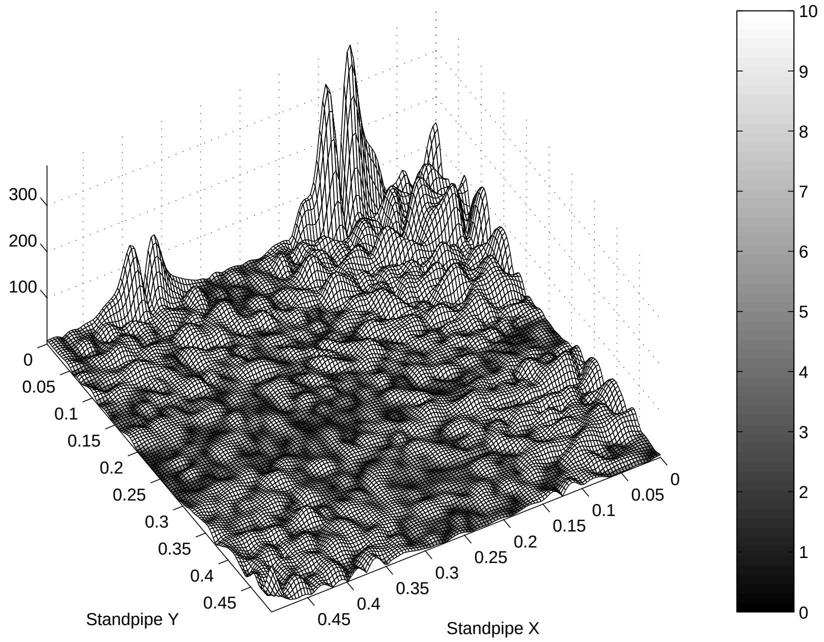


Figure 8: Power spectrum of channel average fuel irradiation calculated using a single taper, and square root transformed to reveal detail.

for the geometry effects. The third term, $h_{ij}$, is the low frequency component and the final term is taken to be independent regression errors.

The development of this model is a staged process, which is not yet complete. For example we still have to introduce a term for the non-central effects of fuel irradiation on surrounding channels. The first stage followed from the correlation and power spectrum analysis previously described. The function $f$ was initially approximated as continuous piece-wise linear, by introducing as separate linear regressors the functions of $x$ defined, for $k = 0, 1, \ldots, 7$ by:

$$s_k(x) = \begin{cases} 0 \\ 1000 + (x - 1000k) \\ 1000 - (x - 1000k) \\ 0 \end{cases} \quad \text{for} \quad \begin{cases} x < 1000k - 1000 \\ 1000k - 1000 < x < 1000k \\ 1000k < x < 1000k + 1000 \\ x > 1000k + 1000 \end{cases} . \tag{2}$$

We have recently replaced this regression by the sum of a linear and exponential response which closely followed the pattern revealed by the piece-wise linear response.

For the geometry effects, some regressors were constructed as patterned occurrences of the values 1, 0 and -1. Some peaks in the spectrum, for which no single such patterned source was obvious, were represented by the pairs of regressors, $\sin(2\pi i/p)$ and $\cos(2\pi i/p)$, with appropriately chosen period $p$, and possibly multiplied by the same functions of the other co-ordinate variable $j$. In total 16 such regressors were found to have a significant effect.

The slowly varying component was modeled by all products of pairs of sinusoidal terms defined as $\sin(\pi ik/n)\sin(\pi jl/n)$, $\sin(\pi ik/n)\cos(\pi jl/n)$, $\cos(\pi ik/n)\sin(\pi jl/n)$

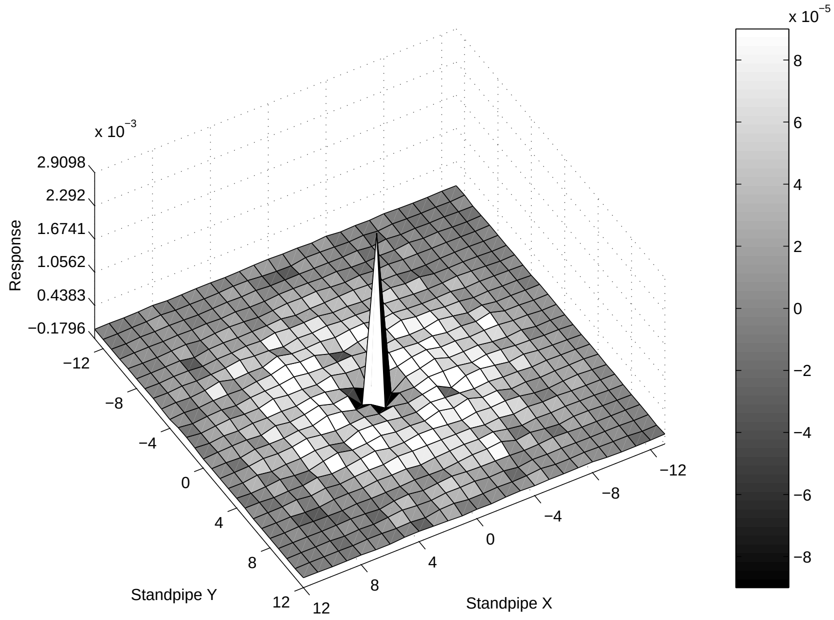


Figure 9: Estimated spatial impulse response of channel average fuel irradiation on CGOTs. This is calculated between the scan temperatures corrected for geometry effects and the non-linearly transformed irradiation.

and $\cos(\pi ik/n)\cos(\pi jl/n)$, with $k, l = 1, 2, \ldots, 10$ and $n = 62$. These functions include both the harmonic and semi-harmonic frequencies; those with a whole number of cycles in the span of 62 pitch units and those with an extra half cycle. Using only the harmonic frequencies would give a periodic fit so that responses at one extreme of the region would affect the fit at the opposite side of the region. These 400 terms were fitted with random coefficients specified to have zero means and variances $d_{kl}^2\sigma^2$, where $d_{kl} = (k^2+l^2)^{-\phi}/\gamma$, $\phi$ is a fractal power, set to 1.25, and $\sigma^2$ is the variance of the model error term. The fractal power was identified by inspection of a log-log plot of the power spectrum of the scan, that had been corrected for the geometry and irradiation effects.

Such a specification is equivalent to assuming a stationary isotropic covariance structure for the slowly varying component. The number of regressors is sufficiently large to provide good local approximation, but still allows the model to be fitted to nearly 4000 responses in a relatively short time. The random coefficients are fitted along with the fixed effects due to fuel irradiation and reactor geometry, in the standard mixed model formulation, see Snedecor and Cochrane (1989). This is the same as for standard regression, but with constraints on the size of the random coefficients, determined by the variance formula. These constraints reduce the effective number of random coefficients and so prevent over-fitting of the model. The scale parameter $\gamma$ in this formula controls these constraints, and is selected to mimimise the mean square prediction error of the model.

Two forms of prediction error were assessed. The first was the cross-validation error, $p_{ij}$. This is the error of prediction of $y_{ij}$ when that response is not used to fit the model. It is generally slightly larger than the residual $\hat{e}_{ij}$ obtained when the model is fitted to all the responses, and its use also guards against over-fitting the model. We may speedily calculate $p_{ij} = \hat{e}_{ij}/(1 - \ell_{ij})$, where $\ell_{ij}$ is the corresponding leverage from the regression. The sum of the leverage values over the responses indicates the effective total number of fixed and random coefficients in the model. The second form of prediction error is that obtained when the model is fitted using only the CGOTs on a $3 \times 3$ sub-grid, then used to predict the remaining responses. This corresponds directly to the main predictive application of the model, and is more robust to model mis-specification.

In fact, rather different values of $\gamma$ were determined by using the two kinds of prediction errors, suggesting that there is opportunity for improvement in the specification of the low frequency component. We were content to use the value of $\gamma$ determined by the, more robust, second method, because it did not lead to a substantial increase in the cross-validation sum of squares for the first method. The effective number of parameters used to fit the low frequency component was then close to 150, one for every 25 responses, which corresponds closely to previous results using kernel prediction.

In our development of the model we first introduced some of the main geometric effects as regressors. We then re-examined the power spectrum of the residuals which revealed effects of a smaller magnitude, that were then added to the regression as sinusoidal terms at appropriate frequencies. After consideration of these effects we found that most such terms could be replaced by indicators associated with identified patterns in the reactor geometry.

We also reviewed the impulse response function estimation using cross-spectral analysis between the temperature scan and the fuel irradiation, after correcting the scan for the estimated geometric effects. This improved the impulse response estimation because the regular geometric effects had added to the noise in the spectral analysis. After estimating the piece-wise linear response of the fuel irradiation we constructed a single transformed irradiation variable in the shape of this response, and with this we repeated the impulse response function estimation.

The spatial regression model has therefore been gradually refined and its predictive accuracy improved. Features of the power spectrum of the residuals, and the cross spectrum between the residuals and the transformed irradiation, reveal some small remaining potential for improving the formulation of the response to the various components of the model.

# 5 Results and comparison with the physical model.

Table 1 presents a list of the main effects coefficients and their $t$ values. Many of these are highly significant. The goodness of fit as measured by the root mean square prediction error of the CGOTs, was 2.56°C. This is for the prediction of each CGOT from all the remaining values. Using the coefficients of the main effects determined from this model, the root mean square error of interpolation from the $3 \times 3$ sub-grid of CGOTs is slightly higher at 2.84°C. This is about half of the value obtained by LTW using kernel prediction. Figure 10 shows the interpolated CGOTS obtained in this way. The improvement comes in part from explaining a high proportion of the variability by

Table 1: Main effects of the spatial regression model

| Coefficient | $t$ value | Description |
|---|---|---|
| 0.000317 | 5.14 | Linear irradiation age |
| -14.5617 | -38.19 | Exponential irradiation age |
| -8.949 | -2.18 | Coolant flow measure |
| 2.215 | 27.64 | East-west period 4 sine term |
| 2.091 | 26.16 | East-west period 4 cosine term |
| 1.591 | 25.92 | Graphite brick size alternating pattern |
| 0.595 | 3.27 | Control rod effect on adjacent CGOTs |
| -0.709 | -8.37 | Banded brick size association with interstitial holes |
| 1.252 | 3.01 | Inner region indicator |
| 1.229 | 5.58 | Double-dwell fuel indicator |
| -0.625 | -10.11 | Sine term side-band of period 4 |
| 0.271 | 4.49 | Cosine term side-band of period 4 |
| -0.394 | -4.37 | Period 8 sine-cosine product, term 1 |
| 0.945 | 8.46 | Period 8 sine-cosine product, term 2 |
| -4.58 | -14.46 | Sector control rod insertion indicator |
| -3.33 | -13.18 | Lift control rod insertion indicator |
| 1.711 | 5.19 | Absorber presence - outer region |
| -2.02 | -10.22 | Absorber presence - inner region |

the effects of the reactor geometry and fuel irradiation. Because of its smoothness, the remaining low frequency component may also be interpolated relatively accurately at the points in between the sub-grid. This fit is somewhat better than that obtained from the physical model, which was adjusted using power-shaping with a higher effective number of free parameters. The physical reactor model has been developed by the industry since the 1960s. It's great advantage is that it is applicable over a wide range of reactors, given the specification of the reactor design, materials, fuel types, operating conditions and other features which have physical properties determined beforehand. Some of the detailed effects, for example those of the alternating pattern of graphite bricks sizes, are not however included in the model.

The physical model also requires some form of calibration, to give good predictions of the CGOTs. Some of its parameters may be adjusted to obtain predictions closer to the measured CGOTs. The statistical model is not however limited in the same way by a prescribed structure, and given adequate data might therefore be expected to provide a closer fit. It uses sound principles of statistical inference to justify the terms which account for significant variation in the data, and to describe the remaining random variation. Its empirical nature is well illustrated by our modeling of the geometry effects. There may be many physical means, some of which may not yet have been identified, by which the geometry influences the CGOTs, but in an empirical model these can be lumped into a small number of regression coefficients. The main limitation of the statistical model is that it cannot be transferred from one reactor to another without re-estimation and in some case reconstruction of the model terms.

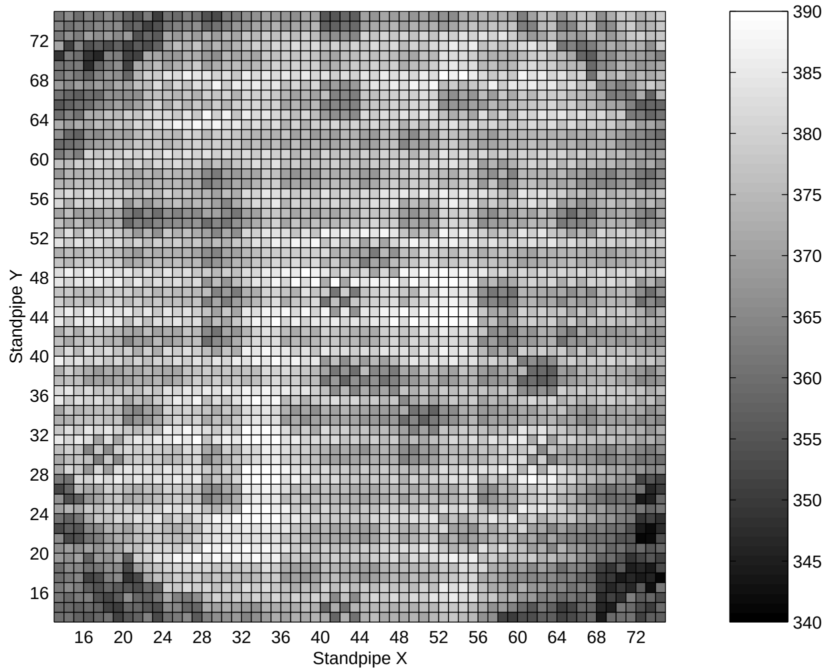


Figure 10: Predicted CGOTs obtained by interpolation from a $3 \times 3$ subgrid using the spatial regression model. These predict the inner square subregion of Figure 1.

Insight from the physical model does influence the statistical model building and some findings of the statistical model may be incorporated into the physical model. The statistical model is however being developed in parallel, providing alternative sets of predictions which will be able to validate conclusions drawn from the physical model.

# 6 Conclusions

We have demonstrated how a range of statistical methods may be used to develop a predictive model of reactor temperatures. The model has been improved in stages and some effects, especially the diffusion effect of fuel irradiation, are yet to be introduced into the model. We have yet to extend the model from the square sub-region to the whole reactor. There are also some aliasing difficulties with identification of the geometry effects using only CGOTs on the $3 \times 3$ subgrid.

The model has, however, already provided more accurate predictions than the physical model, and these may be used to improve the assessment of reactor risk associated with unobserved temperature fluctuations. The quality of the scan data is important to many aspects of the reactor operation and the model has the potential to contribute to this data quality. Scan data may be validated by inspecting residuals from the model predictions and checking that no anomalous values are present. For example, scan data from one standpipe used in our analysis was questioned because the CGOTs were noticeably higher than those in surrounding standpipes. Residual inspection confirmed, however, that these observation could be explained without difficulty by the effects included in the model.

## Acknowledgements

We acknowledge, with thanks, the supply of reactor data by Magnox Electric plc, the use of their physical model, and valuable discussions with their staff at the Berkeley Centre. Financial support from the Industrial Management Committee of the nuclear industries under contract IMC/RPS/5003(a), and from the Engineering and Physical Sciences Research Council under grant GR/M13848, is gratefully acknowledged.

## References

Box, G.E.P. and Jenkins, G.M. (1976). *Time Series Analysis, Forecasting and Control*, Holden-Day.

Hastie, T.J. and Tibshirani, R.J. (1990). *Generalized Additive Models*, Chapman Hall.

Logsdon, J. and Tunnicliffe Wilson, G. (2000). Prediction of extreme temperatures in a reactor using measurements affected by control action. *Technometrics* (under revision).

Scarrott, C.J. and Tunnicliffe Wilson, G. (2000). Spatial Spectral Estimation for Reactor Modeling and Control. Presented to 2000 Joint Research Conference on Statistics in Quality, Industry and Technology, University of Washington. At www.maths.lancs.ac.uk/~scarrott/presentations.html

Slepian, D. (1964). Prolate spheroidal wave functions, Fourier analysis and uncertainty. Part 4: Extensions to many dimensions; generalized prolate spheroidal functions. *Bell Systems Tech. J.*, **43**, 3009-3057.

Snedecor, G.W. and Cochrane, W.G. (1989). *Statistical Methods* (eighth edition). Iowa State University Press, Ames.

Zhang, X. (1994). Wavenumber spectrum of very short waves: an application of two-dimensional Slepian windows to spectral estimation. *J. Atmos. Oceanogr. Tec.*, **11**, 489-505