APPLIED STOCHASTIC MODELS IN BUSINESS AND INDUSTRY
Appl. Stochastic Models Bus. Ind., 2004; **20**:219–238 (DOI: 10.1002/asmb.524)

Spatio-temporal modelling of corrosion in an industrial furnace

John Little[1,*,†], Michael Goldstein[1], Philip Jonathan[2] and Kees den Heijer[3]

[1] *University of Durham, U.K.*
[2] *Shell Research, Chester, U.K.*
[3] *Shell Research, Amsterdam, The Netherlands*

SUMMARY

Optimal inspection and maintenance of complex systems in modern industry is important for safety and economic reasons. With appropriate statistical modelling, the utilization of inspection resources and quality of inferences can be greatly improved. Modelling and inspection of a full-scale industrial furnace subject to corrosion will be considered. A suitable Bayesian spatio-temporal dynamic linear model is developed for wall thickness, by eliciting the beliefs of experts and incorporating other relevant data for related systems. The model may be used to derive efficient inspection schedules for corrosion detection and we demonstrate the considerable reduction in the inspection burden which the model allows. Copyright © 2004 John Wiley & Sons, Ltd.

KEY WORDS: spatio-temporal; DLM; corrosion; inspection; correlation; simulation

1. INTRODUCTION

The integrity of many industrial systems has to be monitored. Often, when expensive non-intrusive inspections have to be made and critical decisions taken as a result, measurements are made at isolated points and used to check the integrity of the entire system. A suitable model for these intermediate points is essential and allows inferences to be made over the entire structure, as well as allowing us to borrow strength from neighbouring points when short of data.

1.1. Background

The use of statistical modelling for estimation when dealing with wall thickness measurements in particular is historically rather limited. For localized corrosion, however, the application of extreme value analysis is relatively common. Industry guidelines, for example those issued by UK Health and Safety Executive [1, 2] treat the estimation of corrosion rate very generally.

*Correspondence to: John Little, Department of Mathematical Sciences, University of Durham, South Road, Durham DH1 3LE, U.K.
† E-mail: john.little@dunelm.org.uk

Contract/grant sponsor: EPSRC CASE studentship award

These guidelines warn of the problems of changing parameters over time, but do not give any specific advice on modelling. Bayes theorem has been used for many years to update beliefs using inconclusive observations by incorporating expert opinion and historical evidence.

There is an extensive spatio-temporal literature often motivated by environmental problems which are commonly temporally rich in data. A common approach in these circumstances (e.g. References [4, 5]) is to use a dynamic linear model (DLM) framework to update uncertainties about general model parameters as observations are made.

Kriging is a common method used for making spatial inferences. The use of Kriging in spatial temporal problems is demonstrated by Mardia *et al.* [6] and Huang and Cressie [7]. In Reference [5] a DLM is used directly in time but only a restricted number of spatial points are modelled. Inferences for intermediate points are then made using Kriging. A particular advantage of this approach is that data is not restricted to a lattice framework and the updates have a relatively low computational load. However, accurate modelling of localized spatial variation is likely to be more difficult.

Many authors have considered ways of reducing the computational load of the DLM. A common approach is to reduce the number of dimensions for instance by constructing the update using summary variables. This sort of technique is most appropriate for systems which are relatively smooth across space, that is to say strongly spatially correlated. Of particular interest is Reference [8]. The reduction of the number of inspections made within a spatio-temporal problem is considered in Reference [9] based on the model constructed in Reference [10], where inspection locations are fixed.

1.2. Outline of paper

In Section 2, we introduce the motivating example—an industrial furnace used in the oil refining process—and discuss the limitations of the existing model. We should however emphasize that this paper is about the general role of spatio-temporal models in constructing efficient inspection schemes, which is of much wider application than the motivating example. In Section 3, we consider a spatio-temporal dynamic linear model which can be used to model continuous systems where measurements are made over time on a finite grid, and discuss the covariance structure used in our particular spatial model. In Section 4, the data set that we have introduced is discussed. Prior beliefs about the means, variances and covariances for this data set and their assessment via expert elicitation are discussed in Section 5. In Section 6, the model is validated using a range of diagnostics and a sensitivity analysis to assess model robustness to changes in beliefs. In Section 7, we discuss future inspection plans, when and where inspections should be made, and different criteria for making these choices. We do this by examining a range of possible inspection strategies and comparing their expected outcomes. Making analytical inferences about our confidence in the overall system is hard because our beliefs about different parts of the system are dependent. Therefore, in Section 8, simulations are used to create realizations of the updated beliefs to allow assessment of the overall risk of the system and to compare directly the best of the candidate inspection schemes from the previous section. The considerable improvements that can be made in the case of the example problem are also discussed. In Section 9, various issues arising from the approach are discussed.

2. MOTIVATING EXAMPLE: CORROSION IN AN INDUSTRIAL FURNACE

Corrosion occurs in many situations and it is often important to know the extent of the corrosion at any particular location. Early identification of corrosion can help in the initiation of preventative measures and accurate long-term prediction is important for forecasting the remaining life of the system. Within the oil industry corrosion affects furnaces, pipelines, storage tanks, valves, nozzles and many other systems. The model presented is potentially suitable for all these applications.

The current lack of a suitable spatio-temporal model within the industry leads to the undervaluing of collected data. The result is that collected data is often tested against a critical value but not actually recorded for future reference. This gives a very limited choice of data sets for retrospective analysis. One of the objectives of this paper is to show the advantages of full recording and analysis of all measurements that are made.

The example described concerns the modelling of corrosion in a full scale industrial furnace. The furnace is located in the lubricating oil production facility at a large European oil refinery. The lubricating oil production facility takes long residue from the refinery's crude distiller unit, separating and purifying them, via a series of distillation columns, de-asphalting and de-waxing units, into lubricating oils with specific viscosities, such as spindle oil and machine oils. Following initial separation of lighter fractions from the long residue, the furnace under consideration is used to heat the remaining hydrocarbons prior to further separation. An outline description of the furnace follows:

This furnace has four sides (N,E,S,W) with one coil of eight pipes on each side (see Figure 1). For each coil (identified by the first number in Figure 1) the hydrocarbon ascends tube 1, descends tube 2, ..., descends tube 8. In Figure 1, entrances to the coils are located in the NW and SE corners (identified by the arrows). Each pipe is 21 feet (6.3 m) tall.

Inspection of the furnace involves major disruption, if not complete shut-down, of the production process at considerable economic cost. For this reason, inspections tend to coincide with full shut-downs of the production facility for major scheduled maintenance and improvements. Typically, such shut-downs occur once every 3–5 years. The main issues for inspection of the furnace are the detection of corrosion in the heating tubes, and estimation of the rate of corrosion. A related concern is optimizing the value of inspection resources, balancing inspection intensity against precision of inferences.

The data is presently modelled with a multiprocess dynamic linear model based on Reference [11]. This model allows forecasting of future corrosion at measurement sites but has no spatial element and so (i) requires data to be measured at the same locations for every inspection, and (ii) gives no information about locations intermediate to the measurement points. The multiprocess model naturally flags outliers and substantial changes in level or slope. Such a multiprocess model requires many parameters to be specified *a priori*, and re-estimated each time data is collected. It is now felt by experts using the multiprocess model, that for most applications such a model is unnecessarily over parameterized given the limited data. It is also felt that a spatial model coupled with diagnostic warnings to identify unexpected observations and manual interventions to recalibrate model parameters when necessary, would better represent the process.

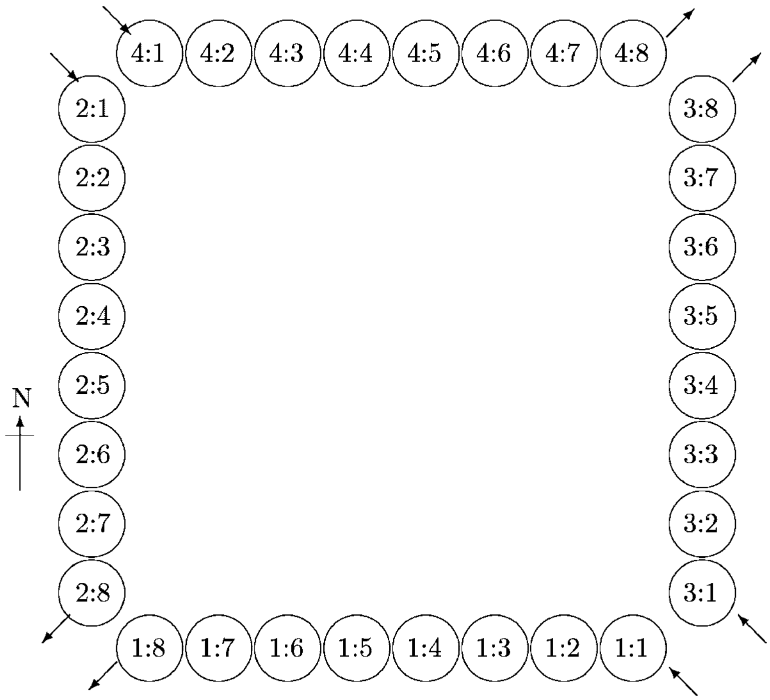

Figure 1. Descriptive diagram of the furnace.

3. SPATIO-TEMPORAL DYNAMIC LINEAR MODELS

A Bayesian spatio-temporal model is a natural way to model many industrial systems with a continuous spatial element that develops in time, where there is extensive experience on the part of the operators but a relatively small amount of data. We develop the approach using a multivariate spatio-temporal DLM. The DLM is developed extensively in Reference [12]. We follow their notation when appropriate.

The motivating example, in common with many industrial systems in the petrol-chemical industry, is expected to suffer from wall loss over the course of its life. Any model of the system should reflect this general trend. Experts consider that it is reasonable to assume that the system will on average deteriorate continuously if operating conditions remain constant, and that current corrosion rate is the best estimator of the future corrosion rate. Therefore, the corrosion is modelled with a locally linear trend.

The system is modelled with three coupled equations: an observation equation, a system equation and a system slope equation.

$$y_{ts} = \theta_{ts} + v_{ts} \quad v_{ts} \sim \langle 0, V_s \rangle$$
$$\theta_{ts} = \theta_{(t-1)s} + \beta_{ts} + \omega_{\theta ts} \quad \omega_{\theta ts} \sim \langle 0, W_{\theta s} \rangle$$
$$\beta_{ts} = \beta_{(t-1)s} + \omega_{\beta ts} \quad \omega_{\beta ts} \sim \langle 0, W_{\beta s} \rangle$$

where y_{ts} is the observed wall thickness value, θ_{ts} is the actual wall thickness and β_{ts} is the corrosion rate all at time t and location s . In this example V_s, $W_{\theta s}$ and $W_{\beta s}$ are constant in time but dependance on t can be introduced if it is considered appropriate. It is further assumed that given V_s, v_{ts} is correlated over s for a given t but is independently and identically distributed (IID) for different values of t. Similarly, given $W_{\theta s}$, $\omega_{\theta ts}$ is correlated over s, but IID over t and, given $W_{\beta s}$, $\omega_{\beta ts}$ is correlated over s, but IID over t. Prior uncertainty about θ_0 and β_0 is similarly captured in $C_{\theta 0}$ and $C_{\beta 0}$ respectively. Given V, W_θ, W_β, $C_{\theta 0}$ and $C_{\beta 0}$, v_t, $\omega_{\theta t}$, $\omega_{\beta t}$, θ_0 and β_0 are all mutually independent. v_{ts} represents observation and calibration errors which will not have any effect on the actual wall thickness, $\omega_{\theta ts}$ represents shocks in wall thickness, for example due to mishaps in operation and $\omega_{\beta ts}$ represents changes in corrosion rate, for example due to long term changes in operating conditions.

V, W_θ and W_β can either be considered known, and then $\omega_{\theta ts}$, $\omega_{\beta ts}$ and v_{ts} can be modelled using the Gaussian distribution, or the Bayes linear approach may be taken, considering V, W_θ and W_β as prior variances for the corresponding error terms. For general information about the Bayes linear approach, see Reference [13], and in the context of DLM analysis, see Reference [12].

The system update variance matrix of a DLM may be constructed from a sum of suitable variances and covariances. This maintains the positive definiteness required of covariance matrices. In the example, the covariance structure between two points is based on the difference in their height, the difference in their pipe number, the difference in their coil number and the difference in distance through the coil (as the hydrocarbon flows). In the example

$$\mathrm{cov}(\omega_{its}, \omega_{its'}) = W_{itss'} = \sum_{j=1}^{4} \Theta_{ij} \exp^{-c_j |d_j|_i^{\Phi}} \qquad i = 1, 2 \tag{1}$$

where Θ_{ij} is a measure of the importance of the component, c_j is a relative scaling parameter (> 0), $\Phi_i \in [1, 2]$ is a general constant and d_j is a Euclidian measure of distance (in some metric) between s and s' for the j^{th} component of variation. Alternative functions are possible—see for example Reference [14].

3.1. Analysis of the model

In the example, the model has monthly time steps, although data is only observed, irregularly, every 3–5 years, and measurement times were recorded to the nearest month. The months where no data was observed are treated as missing values. Missing data is naturally handled within the DLM framework. We have already assumed either a Gaussian distribution where means, variances and covariances for the errors are taken as known or equivalently a Bayes linear framework where we have estimates for the means, variances and covariances of the errors. The model allows means, variances and covariances between wall thicknesses to be calculated at all time points and spatial locations. So given any subset of wall thickness observations beliefs about all the means, variances and covariances for every θ_{ts}, β_{ts} and y_{ts} across the entire system can be updated either by Gaussian conditioning or Bayes linear adjustment. Following observations being made, at any collection of locations and time points, forecast means and variances for any locations now and in the future are then calculated directly from the system equations based on the updated beliefs for all parameters at the current time.

4. DATA

In this section data from the furnace described in Section 2 are discussed.

Vessel wall thickness was measured at heights 3 feet (0.9 m), 7 feet (2.1 m), 12 feet (3.6 m) and 18 feet (5.4 m) on every pipe in August 1983, October 1985, November 1986, March 1988, June 1990, July 1993 and September 1996. Henceforth, the units of height are in feet (units of 0.3 m). Approximately 10% of the data are missing—mainly at two particular points in time. We consider the furnace to have come on line at the start of 1982 (time 0) so the first wall thickness data were recorded 20 months later. All data are recorded to one decimal place except at time 46 when it is rounded to the nearest $\frac{1}{2}$ mm. The data are ultrasonic thickness readings, made using a compression probe and spot readings. It seems likely that some of the observations actually reflect a transient thickness possibly caused by interference when measurements are made. In particular the change in corrosion trend which occurred around 59 months shown in Figure 2 cannot be attributed to any particular factor by the experts. Following common practice the transient thickness is used as a surrogate for true wall thicknesses where this additional error will form a large part of the system uncertainty.

Visual inspection of the data reveals one outlier corresponding to coil 3, tube 2, and height 12 at time 75 when a measurement of 1.4 mm was recorded. This is clearly a recording error of some description. This value is treated as missing which then fits in with the rest of the data at time 75 when only heights 3 and 7 were measured.

Initial mean wall thickness is 11.39 mm and final mean wall thickness is 10.96 mm. It can be seen in Figure 3 that the wall thicknesses at different locations are correlated. In this case there may be particular concerns about the pipe above the height 18.

Figure 4 shows every measurement at height 7 in coil 1. Wall thicknesses for different pipes generally increase and decrease simultaneously, for example, at time 20 pipes ordered by wall thickness are 1, 2, 3 & 7, 8, 4 & 6 and 5, and at the last measurement they are ordered 1, 2, 4 & 7,

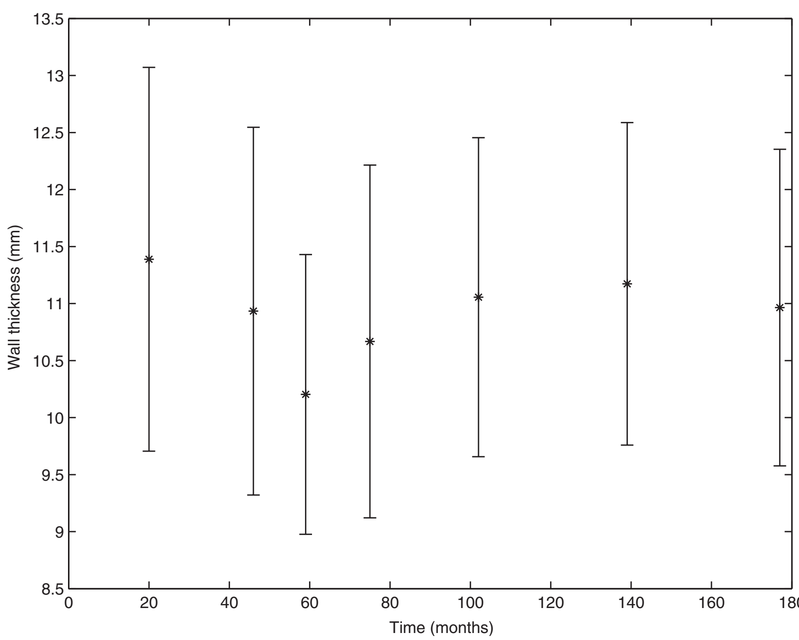

Figure 2. Mean wall thickness ± 3 data S.D.

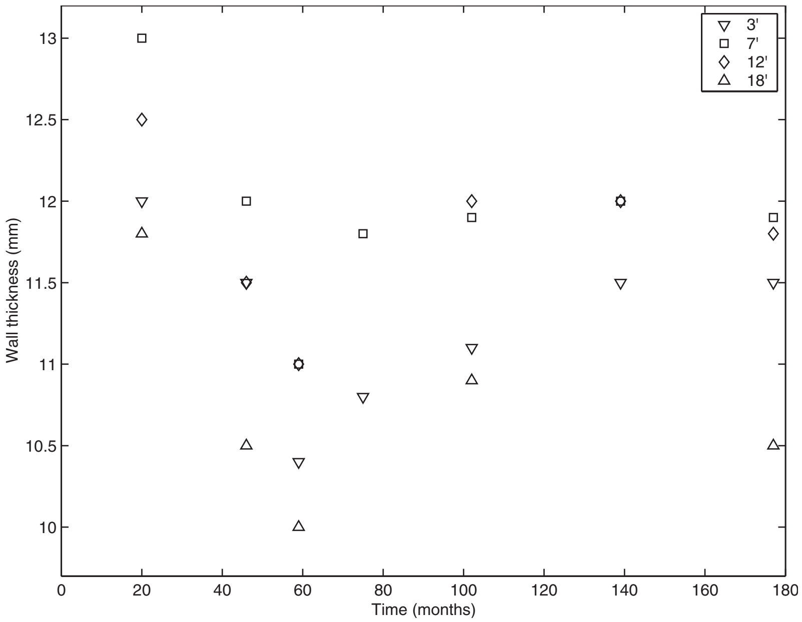

Figure 3. Wall thickness for pipe 1.

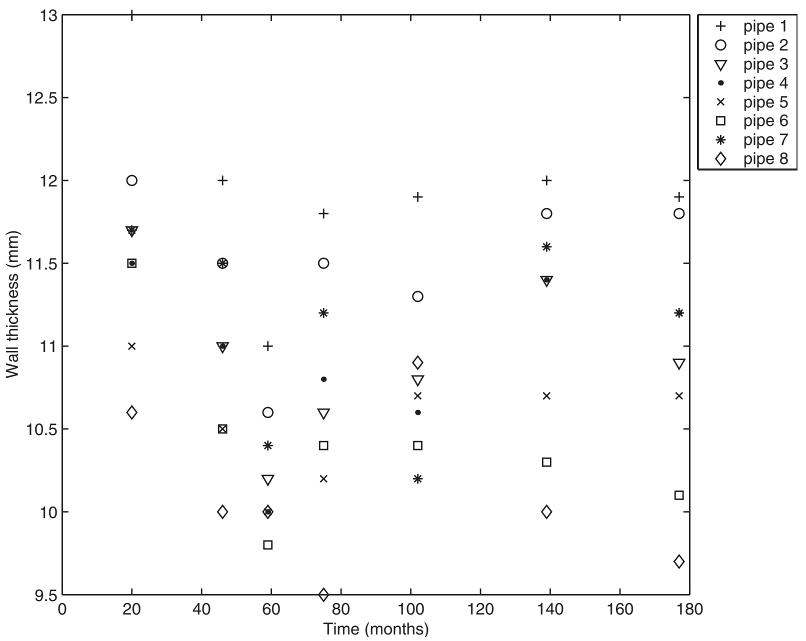

Figure 4. Wall thickness for height 7 across coil 1.

3, 5, 6 and 8, showing less rearrangement than might be expected. As expected, this supports the beliefs of the system experts that remaining wall thicknesses are correlated across pipes. For this particular data set correlations between adjacent pipes are of the order of 0.7.

Experts believe the corrosion will be about 0.3 mm per annum. Initial corrosion is in the range 0.2 mm per annum with a maximum mean corrosion (preceding the inspection at time 59) of about 0.7 mm per annum. Net corrosion over the whole period is approximately 0.025 mm per annum. The standard deviation of the mean corrosion rates is about 0.5 mm, but, within the model, the variance of the mean corrosion rate can be set lower than this value, since the corrosion rate is known to be positive and is believed to be reasonably constant (i.e. this value includes observation noise which is modelled separately).

5. PRIOR BELIEFS

In this section we describe the choice of parameters for the spatio-temporal model used in this example.

5.1. General considerations

Inspection procedures are usually determined by a broad risk based inspection scheme. Risk assessments in the context of inspection are discussed in general in Reference [15]. The inspection scheme will specify the requirements for the inspection tool, the maximum period allowed between inspections (the inspection interval) and the design and implementation of the scheme—examples of these aspects of the inspection scheme may be found in References [16, 17].

For a detailed analysis of a particular system with a particular inspection tool, specific and detailed literature should be consulted (e.g. Reference [18]). The operators of industrial systems, like those studied here, will usually have extensive reports about any given vessel, and the operating system and conditions. Due to the technical nature of many of these reports, they generally need to be interpreted by experts in the field. Our industrial collaborator considers it good practice to base inspection regimes on experience with several vessels of similar nature which have run under similar operating conditions for substantial periods of time (see Reference [16] for further details).

In general, it is recommend that model fitting is based on careful analysis of data collected on such similar systems, in combination with careful elicitation of expert beliefs; the general process of elicitation for complex spatial systems is discussed in detail in Reference [19]. This may raise challenging statistical problems but is unavoidable if reliable and efficient inspection is to be carried out.

5.2. Parameter choice for the example

For the purpose of illustration in the current example, parameter values were chosen consistent with the observed data, which were also considered plausible by expert assessment.

The manufacturing specification of the vessel gives an initial wall thickness of 11.5 mm which is consistent with the observed means. Initial corrosion is set to be $\frac{0.3}{12}$ per month. System experts usually have very strong beliefs that the observation errors are uncorrelated. The observed data for this particular example especially in relation to the 59$^{\text{th}}$ month appears to go against this belief. Modelling correlation within the observational errors may allow more accurate representation of this particular data set. However, we have tried this and it does not appear to make much difference to the forecasts and for the purposes of this example proceed assuming the observation errors are indeed uncorrelated.

Observation noise can be estimated relatively accurately using repeated measurements at one location while system variance is often harder to estimate. Following consultation with experts the observation SD (noise) is believed to be about 0.35.

Current practice suggests that the designer of the existing inspection regime believed it sufficient to measure four locations on each pipe. This may be used as a crude diagnostic check: that is, the analysis should agree that about four points per pipe is an appropriate number of inspections.

A covariance matrix is constructed for the system levels (W_θ) and then the same structure (with a re-scaled variance) is used in the covariance of the system slope (W_β), and in specifying prior level and slope variance ($C_{\theta 0}$ and $C_{\beta 0}$). If more detailed belief structures become available for W_β, $C_{\theta 0}$ and $C_{\beta 0}$ these would be used instead but this seems unlikely. The same correlations are used between points for the update levels, update slopes, prior levels and prior slopes. It is assumed that there is no correlation between these four sets of parameters.

For this illustration parameter values for the model are determined using only expert judgement and data from times 1 to 5. The model is then validated against subsequent times. Data sets exist for other similar furnaces and these could be exploited, in practical applications, to improve the assessment of the covariance structure.

The system update variance is constructed from a sum of suitable covariances, as in (1). Thus the update covariance between any two points on the furnace may be described using:

$$\begin{aligned}\operatorname{cov}(\omega_{its}, \omega_{its'}) = \Upsilon_i(A \exp^{-\alpha|h_s - h_{s'}|^{\Phi_i}} + B \exp^{-\beta|p_s - p_{s'}|^{\Phi_i}} \\ + C \exp^{-\gamma|c_s - c_{s'}|^{\Phi_i}} + D \exp^{-\delta|d_s - d_{s'}|^{\Phi_i}}) \qquad (2)\end{aligned}$$

where h_j is the height, p_j the pipe, c_j the coil, d_j the inside pipe distance around the coil for location j, and all other values are suitable constants. We set $\Phi_i = 1$ since this is considered to be a suitable model which produces a Markov structure which can help reduce computational load. Two points will thus be independent if they are separated by known points.

Eliciting beliefs about spatial covariances is difficult. Experts agreed with the following statement: 'correlation due to two heights differing by k feet' will be two times stronger than 'correlation due to two pipes differing by k pipes', and four times stronger than 'correlation due to two coils differing by k coils' (coil 4 is next to coil 1), but will have the same correlation as 'correlation due to two points as hydrocarbon flows differing by k feet'. This statement is interpreted to give values for $A = 1$, $B = 0.5$, $C = 0.25$ and $D = 1$ in (2).

Following the consultation with experts $\alpha = \beta = \gamma = \delta = 0.003$ and in the overall update variance parameter $\Upsilon_{\text{level}} = 0.00001$ (W_θ) and $\Upsilon_{\text{slope}} = 0.000001$ (W_β) and in the overall initial variance parameter $\Upsilon_{\text{level}} = 0.05$ ($C_{\theta 0}$) and $\Upsilon_{\text{slope}} = 0.005$ ($C_{\beta 0}$).

In both W_β and $C_{\beta 0}$ the overall update value is set to be a tenth of that used in the full model but the correlations are left unchanged.

An alternative to specifying the update uncertainty directly as we have done is to use discounting (see Reference [12]) where the update system variance is modelled as a fixed proportion of the underlying system variance. Discounting requires equally careful prior specification. Discount factors have more effect in longer time series but many industrial inspection problems have very few measurements in time, and discounting is not applied to the model for this reason.

The variance of the double data difference $\mathrm{var}(y_t - 2y_{t-1} + y_{t-2})$ has expectation $2W_\theta + W_\beta + 6V$. With V accurately estimated, if a fixed ratio is assumed between W_θ and W_β, then W_θ and W_β can be estimated directly from the data. In the example because of the short time series and the missing data only a rough estimate can be found and this supports the values discussed above.

A ceiling for the update slope variance can be estimated using the requirement on the update slope to remain negative. Initial corrosion is ≈ 0.02. If at least three standard deviations are required between this value and 0, the standard deviation must be $\leqslant 0.005$ and the variance $\leqslant 2.5 \times 10^{-5}$.

6. VALIDATION OF MODEL

In this section we apply simple checks that the model parameters are suitable including three diagnostic tests and a sensitivity analysis to changes in belief.

A simple check that the model is fitting the data can be made by plotting the system mean for each time and location as well as the data and three standard deviation confidence bounds. Figure 5 is fairly typical of the plots obtained by doing this. Most data fall within the confidence bounds specified for the pipe.

6.1. Diagnostics

Three diagnostic tests are considered at the time of the fifth inspection when fitting the model on the first five data points. These are: univariate forecast error $\underline{y}_t - E(\underline{y}_t \mid D_{\{t-1\}})/\sqrt{\mathrm{var}(\underline{y}_t \mid D_{t-1})}$, which is expected to produce an approximately normal histogram curve, multivariate forecast error $(E(\underline{y_t}) - \underline{y_t})' \, \mathrm{var}(\underline{y_t})^{-1}(E(\underline{y_t}) - \underline{y_t})$ with expectation length(y_t), and data differencing where

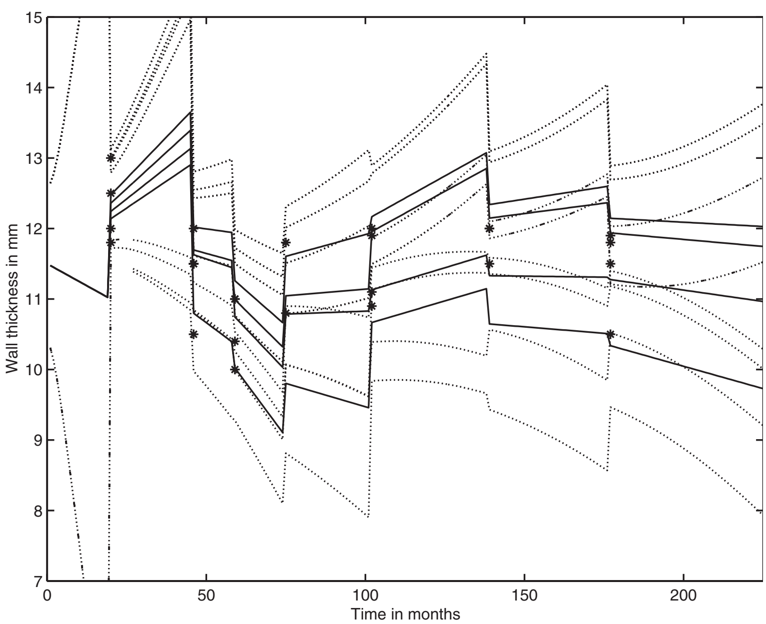

Figure 5. Data and forecasts at the measurement points of pipe 1, coil 1.

$E(\text{var}(\Delta^2 y_{ts})) = 2W_\theta + W_\beta + 6V$, which is applied locally due to the short length of the time series.

Figure 6 shows the univariate forecast error density which shows no cause for concern. The light tails indicate model variances may be slightly too large. The bias of the mean might be attributed to the correlation between the points effectively reducing the sample size. The model correlation structure is exploited to remove the correlation between the points in the multivariate forecast error, where there is an expected value of 640 and an observed value of 780 suggesting prior variances may be slightly too small. The differences between the expected data difference and the actual data difference are in the range $[-1.5, 3.1]$ with a mean of 1.6. The same diagnostics can be applied following the final inspection in September 1996. Figure 7 shows the univariate forecast error, the multivariate forecast error has an observed value of approximately 1061 with an expected value of 1024 and data differencing has range $[-2.6, 2.0]$ with mean 0.0. As before these diagnostics do not indicate any cause for concern.

6.2. Sensitivity analysis

The sensitivity of the model to changes in prior beliefs and update variances can be important. The extent to which the data can distinguish between changes in beliefs may be examined by changing these parameters and then re-running the diagnostics. See Table I.

Univariate forecast error histograms give no cause for concern. Data differencing is fairly consistent and does not raise any diagnostic concerns for the changes considered. Looking at the multivariate forecast error, the observation variance is the most important of the variables to elicit accurately, fortunately it can also be estimated relatively easily using repeated measurements. System correlation is probably the hardest component to elicit beliefs about but since it is fairly robust under all three diagnostics this is not too serious a problem.

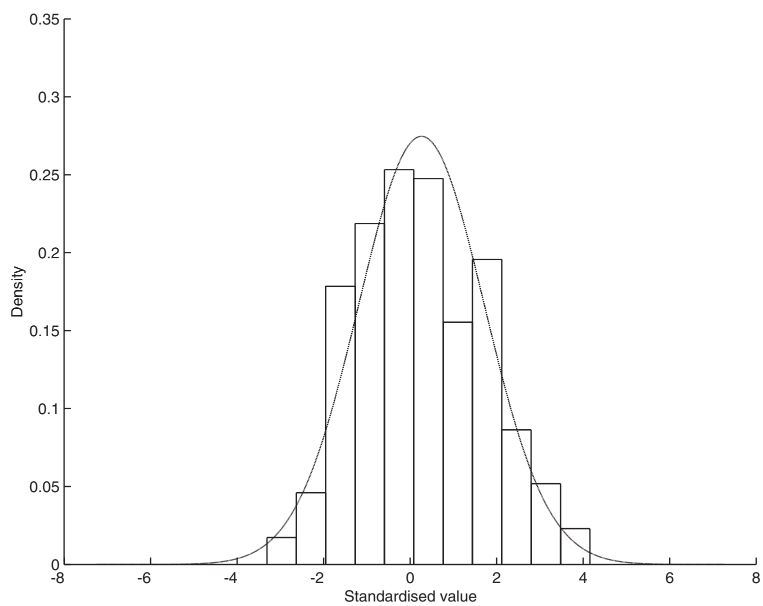

Figure 6. Univariate forecast error density time 5.

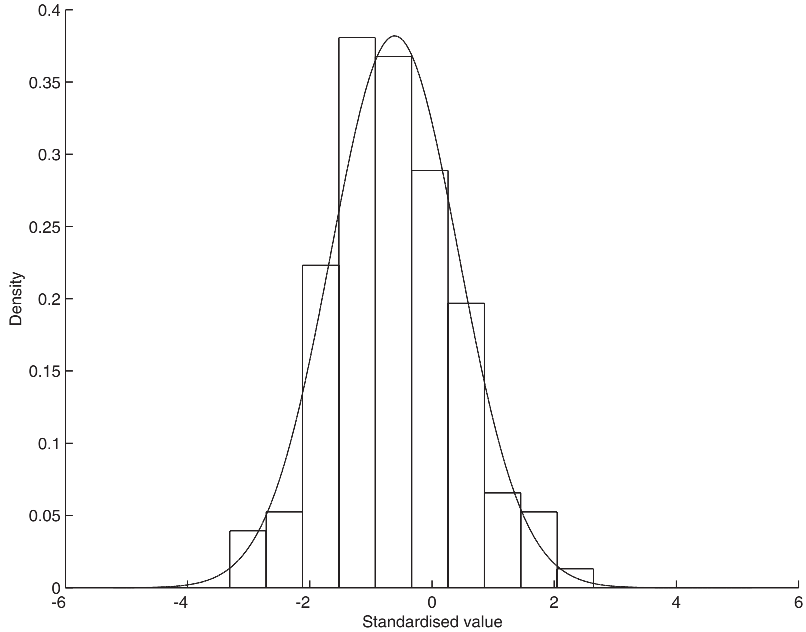

Figure 7. One step ahead forecast error density time 7.

Table I. Changing model parameters.

Description	Data differencing	Multivariate forecast error (Expected = 1024)
Original	$[-2.6, 2.0]$, mean = 0.0	1061.2
↑ (+40%) observation variance V	$[-2.5, 2.0]$, mean = 0.0	562.3
↓ (−40%) observation variance V	$[-2.9, 2.0]$, mean = 0.0	1928.3
↑ (×3) system variance W_θ	$[-2.8, 2.2]$, mean = 0.2	1744.1
↓ (/3) system variance W_θ	$[-2.4, 2.2]$, mean = 0.2	965.8
↑ (×3) system correlation α, β, γ & δ	$[-2.6, 2.0]$, mean = 0.0	1012.2
↓ (/3) system correlation α, β, γ & δ	$[-2.6, 2.0]$, mean = 0.0	1106.4

7. FUTURE INSPECTIONS

In this section some potential inspection schemes and several ways of comparing them are discussed. The change in cost when the number of inspections changes is also considered. We demonstrate for this example the considerable benefits of being able to inspect anywhere rather than using a fixed pattern across the entire system.

Historically, inspections have been made at fixed locations. We consider whether inspecting at alternative locations can help increase efficiency. Using the spatial model, the effect of inspection at any set of points on the system can be analysed. Variance at a point will be lower if that point has been inspected, but the variance will also be influenced by the location of other inspections (both current and previous). If the uncertainty about a point is too high either that point or a collection of points with sufficiently high cross-correlation should be inspected. A formal method is required for deciding where future inspections should be made. This is done by setting one or more constraints and then optimizing. There are many possibilities for these constraints and a few intuitive choices are considered.

For each point of the system (s) there is a minimum required wall thickness c_s ($\geqslant 0$) which the system level is expected to stay above. The inspection at time t should balance the probability of the wall thickness falling below this minimum required value, the inspection cost and the inspection interval k (i.e. the first time when following the inspection the uncertainties become unacceptably large). The inspection interval should be as large as possible, but often it will be overruled by some other constraint, for example, the next scheduled shut down or safety legislation.

Two scenarios requiring attention are high uncertainty and low remnant life. Increased inspection will help the former but do little for the latter. Observations do not change the true life of the system, only the uncertainties about it. The aim is to maximize confidence about the expected life. As a multivariate calculation this is most easily handled using simulation. Ideally for every possible inspection regime a nested simulation would be run where, following an inspection of this simulated data, further simulations would be run to ascertain the maximum risk prior to the next inspection. The most efficient scheme with sufficiently low risk would then be used for the true inspection.

A useful inspection will reduce the uncertainty at and around the inspection location. This reduction in variance may be used alongside a heuristic argument about how near the wall thickness is to the minimum allowed value to avoid the high computational load associated with the simulation approach described above. Two variance and two heuristic criteria are used as a first approximation, see Table II. These criteria cover many intuitive possibilities and can be maximized over possible sampling schemes. Therefore, these criteria are used to identify designs with good heuristic properties and full simulation strategies are reserved for confirming that the chosen design does offer an acceptable basis for sampling inspection.

Before an inspection is made at time t the effect this inspection will have on the uncertainty at time $t+k$ should be considered. Letting s index all spatial locations across the system a heuristic criteria is defined as $h(s,t,k) = E(\theta_{\{s(t+k)\}}\,|\,D_{\{t-1\}}) - c_s/\sqrt{\mathrm{var}(\theta_{s(t+k)}\,|\,D_{t+k})}$. This heuristic value should be large across the entire system. If it is small at any point it is not possible to be confident that the critical minimum wall thickness will not be passed.

Attention is focused on the heuristic criteria, because the variance criteria while easier to visualize, are best suited when the wall thickness is far from the minimum allowed wall thickness. This, however, is the same as maximising the minimum (or mean) precision of any point. If all wall thicknesses are far from the critical level, significant changes in variance will dominate the precision calculation, so there are few advantages over the standardized risk approach. A combination of all four criteria is also a possibility.

To find the optimal inspection scheme for a given budget, every possible affordable combination of measurement points would have to be considered. In addition to the high

Table II. Some possible criteria.

1	Maximum variance	$\max_s\ (\mathrm{var}(\theta_t \vert D_t))$
2	Mean variance	$\mathrm{mean}_s\ (\mathrm{var}(\theta_t \vert D_t))$
3	Mean heuristic	$\mathrm{mean}_s\ (h(s,t,k))$
4	Minimum heuristic	$\min_s\ (h(s,t,k))$

computational load the end result might be impractical to implement, as the inspector would have to locate every inspection point individually rather than following a simple pattern. Instead the improvements to be had from a range of possible sampling strategies that can be implemented using the spatial model are first considered. One hundred percent inspection will give the largest reduction in uncertainty that can be achieved for any inspection plan, and can thus be used as a bench mark. Schemes may either inspect regularly across the system or, in the case of the risk-based scheme, inspect in accordance with our prior beliefs, see Table III.

The intermediate schemes inspect points on every pipe midway between the existing inspection locations. The alternating option may be further extended by swapping which points are inspected for subsequent inspections. In this example, the risk-based scheme categorizes as high risk any pipes which are believed at some location to be within 3.5 standard deviations of the critical value.

Of the 3 and 4 point schemes (see Table IV) only the risk-based scheme gives much information about the 'minimum heuristic'. 'five intermediate' targets the five major spatial areas of uncertainty between the traditional inspection locations hence 'minimum heuristic' shows a marked improvement over the four point schemes.

7.1. How many measurements

Doing a full search over a given number of inspections per pipe to find the optimal set of inspection locations is computationally very intensive. A stepwise delete (or add) process can be

Table III. Some possible inspection schemes.

Scheme	Description
All	All possible locations
Nothing	No inspection
Existing	Heights 3, 7, 12 and 18 on every pipe
3 intermediate	Heights 1, 10, and 21
4 intermediate	Heights 5, 10, 15 and 20
5 intermediate	Heights 1, 5, 10, 15 and 20
Alternating	Even pipes as 'Existing' and odd pipes as '4 intermediate'
Risk-based	Categorise each pipe as 'low risk' (50%) or 'high risk' (50%) and then apply '3 intermediate' or '5 intermediate' respectively.

Table IV. Different criteria versus different schemes.

Scheme	Maximum variance	Mean variance	Mean heuristic	Minimum heuristic
All	0.53	0.50	13.35	9.97
Nothing	2.30	1.56	4.58	2.64
Existing	2.17	1.33	6.22	2.80
3 intermediate	1.87	1.21	6.21	2.96
4 intermediate	2.11	1.08	6.92	2.73
5 intermediate	1.32	1.00	7.36	4.11
Alternating	2.05	1.08	7.24	2.82
Risk-based	1.55	1.07	6.83	4.10

used instead with a very much reduced load. Starting with full inspection, the inspection at the entire level which is needed least, according to the heuristic, is removed. This process may be repeated until only the required number of heights are being inspected. The alternative stepwise add process may also be used.

Figure 8 shows the expected value of the heuristic against the number of heights inspected. Having a very low number of inspections appears ineffective at reducing the maximum uncertainty and having a few uninspected points dramatically affects results. High confidence can only be achieved by 100% inspection. In this example stepwise deleting gives better results than stepwise adding, and further attention is restricted to this approach. Based on the ordering determined by the minimum heuristic the mean heuristic is also plotted. The slope looks concave but this could be a feature of applying the stepwise delete to the minimum heuristic and not to the mean heuristic. The pipe heights are, from least important, to most important:

7,18,12,3,20,16,9,14,5,11,2,19,17,8,13,6,4,10,15,1,21.

The ends of the pipe are most critical, the previously inspected points least critical. Points adjacent to those still being inspected also have a low inspection priority.

Using stepwise deletions we get four recommended heights 1, 10, 15 and 21 with a heuristic value of 3.31, which is a marked improvement over the original suggestions described at the start of Section 7. Were resources available it would seem sensible to make a fifth inspection at height 4 (or perhaps 5) to exploit the jump in the heuristic value curve. This possibility is investigated later. Further improvements may be achieved by combining stepwise deletions with stepwise additions. However in this case, because of the regular nature of the previous inspections, significant improvement seems unlikely and we proceed with inspections at heights 1, 10, 15 and 21. The risk associated with the set of measurements may be estimated by simulation.

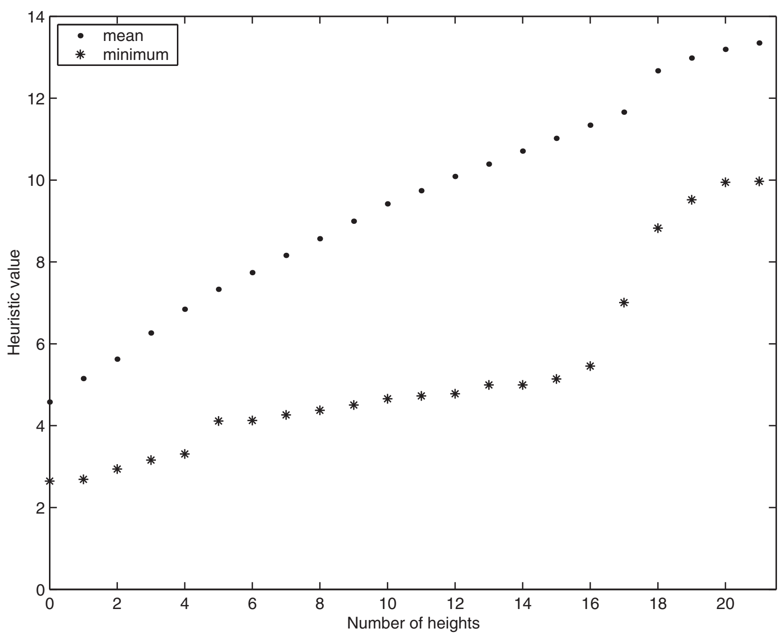

Figure 8. Heuristic value using stepwise delete.

7.2. More advanced inspection regimes

The optimal inspection would comprise of measurements at locations across the entire system and not just at restricted heights across the system. The resulting scheme may be practically difficult to implement and there may be restrictions on where measurements can be made but these concerns are put aside for the present. The expected value of the scheme must be known if the cost of the inspection is to be assessed and compared against alternatives.

An approximation to the optimal scheme, is to sequentially add (or delete) single points according to the criteria. The criteria value is checked at every point and inspected at the worst point. Subsequent points are then selected using the updated beliefs. Where we consider the heights at which inspections are made (see Figure 9) a clustering of points around the areas where prior uncertainty was highest—heights 1, 5, 10, 15 and 21, can clearly be seen. The number of inspections in each coil, in each pipe and generally throughout the system were also considered but no obvious patterns were found. Figure 10 shows the increase in the minimum and mean heuristic values as the number of inspections increases using a sequential add on the minimum heuristic.

The mean heuristic increases continuously with the first inspections being most important. The minimum heuristic increases somewhat sporadically. This may be extenuated as a result of using the stepwise add, but may also be due to a certain intensity of inspection covering entire areas of similar structure.

Comparing inspection based on these 128 inspection locations (i.e. an average of four per pipe) to the best scheme found with fixed inspection heights on every pipe, the minimum heuristic value is much higher (5.11 rather than 3.31) and to get this level of accuracy using

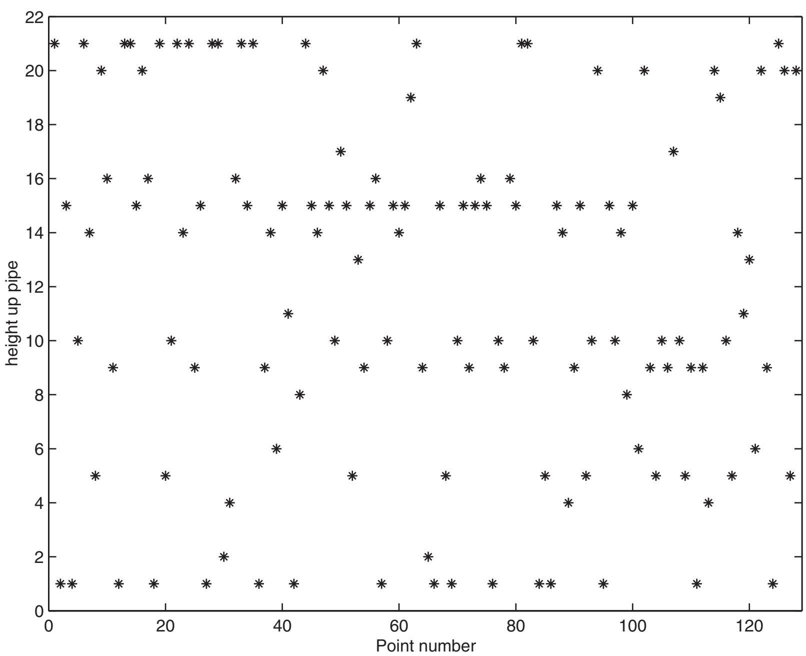

Figure 9. Optimal inspection heights using criteria.

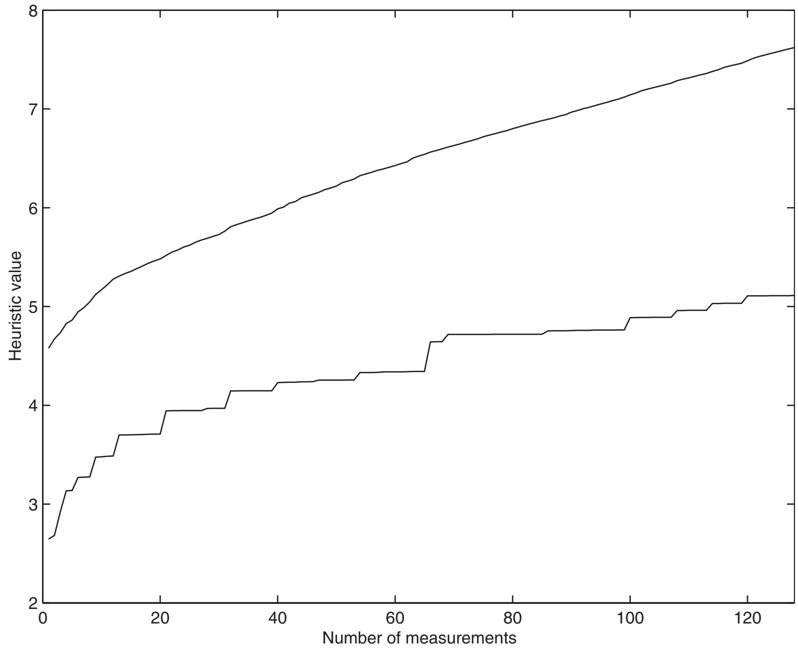

Figure 10. Number of points inspected against heuristic value Upper line—mean heuristic, lower line—minimum heuristic.

standard points 15 inspections would need to be made on every pipe. The mean heuristic is also improved (7.63 rather than 6.85) which would require six regular inspections per pipe. To achieve the same minimum heuristic accuracy as the 128 fixed height inspections only nine carefully positioned inspections need be made.

8. ASSESSING THE RISK OF THE CHOSEN SCHEME

In this section, the risk of system failure in the future is calculated. A demonstration is then made of how simulation can be used to assess how accurately any given inspection regime estimates this risk in comparison with 100% inspection.

The objective is to run the furnace for as long as possible and ideally stop just before wall thickness becomes too thin. Because there is uncertainty about the time of the crisis point a safety margin may have to be included, depending on the consequences of overshooting the minimum wall thickness.

In a multivariate model, calculating the probability that the worst point of the system will pass the crisis point is difficult. By simulating data sets according to the current beliefs for the end point of the forecast, the probability that the wall thickness will fall below the critical minimum value can be found.

One hundred thousand data sets covering the entire system are simulated corresponding to our beliefs at the time of the seventh inspection, given all data up to and including the fifth inspection. In about 7% of cases the minimum wall thickness is below the critical value of 6 mm.

From the simulation study we find the mean wall thickness of the worst point of the system, will be about 8.1 mm with a standard deviation of 0.75 mm. The minimum observed measurement recorded at the time of the seventh inspection was 9.7 mm, which emphasizes the risk associated with the unmonitored intermediate points.

The inspection of all pipes at heights 1, 5, 10, 15 and 21 is compared against the inference that would be made following a full inspection of the same data.

First, given all data from inspections 1 to 5, a series of sets of data are simulated according to the beliefs at the time of the sixth inspection. Each simulated data set is then inspected according to the chosen scheme. The beliefs about what should be expected at the time of the seventh inspection can then be updated. A number of realizations are simulated (300 in the example) of the updated beliefs at the time of the seventh inspection for each data set. The proportion of simulations with a minimum wall thickness below 6mm can then be found. Ideally this proportion will be either 0 or 1 which would allow perfect prediction of the situation at the time of the seventh inspection using the results of the sixth inspection. By comparing with updated beliefs following 100% inspection of the simulated data the reliability of the inspection scheme may be assessed.

The probability of passing the minimum wall thickness for the candidate scheme is plotted against that for full inspection for each realisation of the sixth inspection. Each plot may be split into four (unequal) quadrants at the value of the maximum risk we are prepared to take, e.g. 0.05. Points in the bottom left indicate neither the candidate scheme nor 100% inspection suggest the minimum wall thickness criteria is likely to be reached. Points in the top left quadrant indicate that full inspection suggests the system is safe but the candidate scheme disagrees. Points in the top right quadrant indicate neither inspection regime believes the state of the system to be satisfactory, while points in the bottom right quadrant indicate the full inspection has found a high risk while the candidate scheme has missed it. The cost of a false positive and of a false negative can then be used to value each inspection regime.

In Figure 11 from the 110 simulation there are three points especially likely to give us cause for concern where the five point model claimed a low probability of system failure, while the full inspection identified a high risk. If the threshold is set at a 10% probability of minimum wall thickness passing the critical value, then there are three false positive points and nine correctly identified high risk points.

Study of plots such as this for different designs allows an inspection plan to be identified which balances cost against acceptable levels of risk.

9. DISCUSSION

In this paper, we have described the modelling of the spread of the corrosion, using a spatio-temporal DLM. Such an approach has three main advantages. Firstly, it allows us to incorporate expert judgements into our inspection planning in a natural and powerful way. Secondly, the model allows us the maximum flexibility in the choice of inspection times and locations. Finally, the model is sufficiently tractable to allow us to obtain, at least approximately, optimal inspection procedures, and to obtain, by simulation, an assessment of the risk associated with these procedures.

A natural extension to this paper is to incorporate decision analysis, balancing utilities for the consequences of different forms of unexpected system failure against the cost utility for making

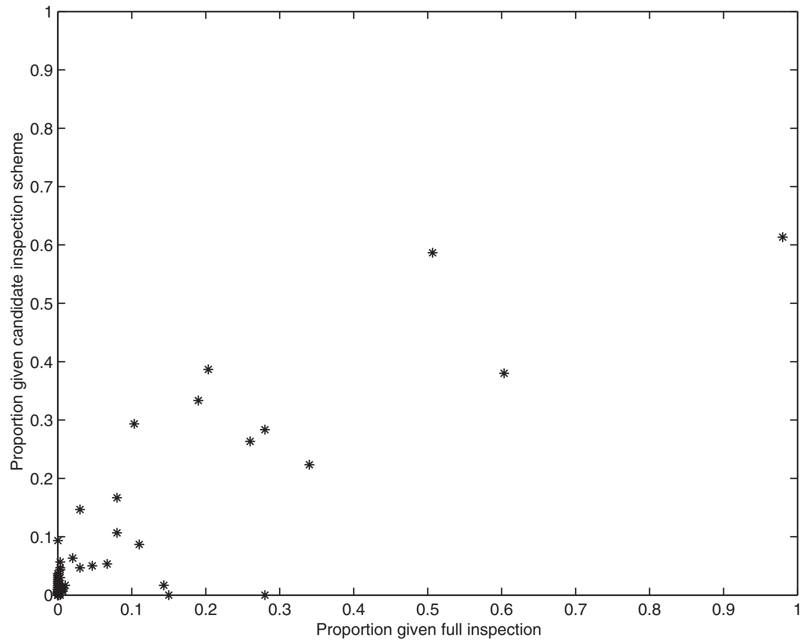

Figure 11. Proportions in risk for full inspection against five point scheme.

inspections. This would allow an assessment of the optimal expenditure on inspection, incorporating an assessment of the cost of the constraints, for example having to inspect when a system is shut down for maintenance.

The model considers changes in corrosion rate as a random walk. This provides a simple and plausible model for the corrosion rates. We have assumed for this example that the corrosion rate change can be modelled with a single variance for all times. However, a changing variance can be handled similarly, for example if suggested by assessment of related data sets. Sudden changes in corrosion rate, as might be experienced if there is a significant change in operating conditions, may be best handled with an intervention. In systems where there is shielding of some description this will reduce the wall loss while the shielding remains intact. A two process multivariate model would be more realistic in these circumstances where we would model a switch from one corrosion rate to the other.

We have assumed fixed values for all our variances and hyper-parameters, a good model fit relying on expert judgement, supplementary data on similar systems and our diagnostic checks. From a fully Bayesian perspective these values will also be unknown and have uncertainties attached to them. This can be especially relevant in time series models over long periods where the variances are more likely to change over the course of the model. Modelling the problem with these additional uncertainties would require the use of more sophisticated simulation models. MCMC methods for DLMs are described in Reference [12]. Uncertainty in the hyperparameters is considered in a space time context by Sansó and Guenni [20] and Wikle *et al.* [21], but both comment on the high computational load of their approaches. Identifying optimal inspection regimes, as in Section 7, using such specifications as these, would be extremely computationally demanding. A simple approach would be to model the uncertainty in the hyper-parameters while updating beliefs given all existing data but to assume these values set at their current expected values while considering future inspection schemes, using the methodology that we have described. A further possibility would be to use the approach that we have described, for fixed choices of the hyperparameters, to restrict the number of possible inspection plans that need to be considered. These plans could then be compared within a fully Bayesian framework, for example using the simulation approach described in Reference [22].

ACKNOWLEDGEMENT

This work was supported under an EPSRC CASE studentship award in conjunction with Shell Research.

REFERENCES

1. UK Health And Safety Executive. Guidelines for use of statistics for analysis of sample inspection of corrosion, 2002.
2. ASTM Standard G16. Standard guide for applying statistics to analysis of corrosion data.
3. Shaddick G, Wakefield J. Modelling daily multivariate pollutant data at multiple sites. *Applied Statistics* 2002; **51**(3):351–372.
4. Stroud JR, Muller P, Sanso B. Dynamic model for spatiotemporal data. *Journal of the Royal Statistical Society* 2001; **63**:673–689.
5. Mardia KV, Goodall C, Redfern EJ, Alsonso FJ. The kriged Kalman filter. *Test* 1998; **7**:217–285.
6. Huang H-C, Cressie N. Spatio-temporal prediction of snow water equivalent using the Kalman filter. *Computational Statistics and Data Analysis* 1995; **22**:159–175.
7. Wilke CK, Cressie N. A dimension-reduced approach to space-time Kalman filtering. *Biometrica* 1999; **86**:815–829.
8. Oehlert GW. Shrinking a wet deposition network. *Atmospheric Environment* 1996; **30**(8):1347–1357.
9. Oehlert GW. Regional trends in sulfate wet deposition. *Journal of the American Statistical Association* 1993; **88**: 390–399.
10. Smith AFM, West M. Monitoring renal transplants: an application of the multiprocess Kalman filter. *Biometrics* 1983; **39**:867–878.
11. West M, Harrison J. *Bayesian Forecasting and Dynamic Models*. Springer: Berlin, 1996.
12. Goldstein M. Bayes linear analysis. In *Encyclopaedia of Statistical Sciences*, Kotz S, Read CB, Banks DL (eds), update vol. 3. Wiley: New York, 1999; 29–34.
13. Cressie N, Huang H-C. Classes of nonseparable, spatio-temporal stationary covariance functions. *Journal of the American Statistical Association* 1999; **94**(448):1330–1340.
14. Dunn MR. Guideline for the use of risk assessment methods. *Technical Report*, Shell, 1997.
15. Terpstra S, Ritchie D, den Heijer C. Non intrusive inspection manual applied to topside facilities. *Technical Report*, Shell, 2002.
16. Terpstra S. Shell expro—Definition of S-NII Model. *Technical Report*, Shell, 2002.
17. Gabrielli C, Huet F, Keddam M, Oltra R. A review of the probabilistic aspects of localised corrosion. *Science* 1990; **4**:266–278.
18. Craig PS, Goldstein M, Seheult AH, Smith JA. Constructing partial prior specification for models of complex physical systems. *The Statistician* 1998; **47**(1):37–53.
19. Sansó B, Guenni L. A nonstationary multisite model for rainfall. *Journal of the American Statistical Association* 2000; **95**(452):1089–1100.
20. Wikle KC, Milliff RF, Nychka D, Berliner LM. Spatiotemporal hierarchical Bayesian modeling: tropical ocean surface winds. *Journal of the American Statistical Association* 2001; **96**(454):382–397.
21. Muller P. Simulation based optimal design. *Bayesian Statistics* 1998; **6**:459–474.

