Statistical Modelling 2004; **4**: 299–313

Efficient Bayesian sampling inspection for industrial processes based on transformed spatio-temporal data

John Little[1], Michael Goldstein[1] and Philip Jonathan[2]
[1]University of Durham, Durham, UK
[2]Shell Research, Chester, UK

Abstract: Efficient inspection and maintenance of complex industrial systems, subject to degradation effects such as corrosion, are important for safety and economic reasons. With appropriate statistical modelling, the utilization of inspection resources and the quality of inferences can be greatly improved. We develop a suitable Bayesian spatio-temporal dynamic linear model for problems such as wall thickness monitoring. We are concerned with problems where the inspection method used collects transformed data, for example minimum regional remaining wall thicknesses. We describe how the model may be used to derive efficient inspection schedules by identifying when, where and how much inspection should be made in the future.

Key words: Bayesian; corrosion; decision support; DLM; industrial statistics; inspection; minima; optimal experimental design; spatio-temporal

Software available from: http://stat.uibk.ac.at/SMIJ
Received September 2003; **revised** August 2004; **accepted** August 2004

1 Introduction

Many large industrial systems including tanks, pipes, vessels and furnaces corrode. It is important to predict accurately the state of these systems for repair, replacement and safety reasons. Inspections are often expensive and the accuracy of the forecasts must be balanced against the cost of making additional inspections. Most inspection methods widely applied in industry record only summary statistics for the area inspected rather than full spatial information. Additionally, these inspections often only cover a small proportion of a system. Common summaries are regional minima or the proportion of the area that has passed a given threshold (for instance, a protective coating). These functional outputs may be due to the design of the data collection device, the time taken to record more extensive details or the requirements placed on the contractor.

Current protocol often checks the observations made against a predetermined criterion (for example, minimum allowed wall thickness) and declares the system to

Address for correspondence: John Little, Department of Mathematical Sciences, University of Durham, Science Laboratories, South Road, Durham DH1 3LE, UK. E-mail: john.little@dunelm.org.uk

 10.1191/1471082X04st081oa

pass or fail accordingly. These data values are then not stored and no direct consideration is given to i) the areas not inspected and ii) the current corrosion rate. Historically the use of statistical modelling for dealing with wall thickness measurements in particular has been rather limited. The application of extreme value analysis for localized corrosion, however, is more common. The estimation of corrosion rate is treated very generally in industry guidelines, for example those issued by UK Health and Safety Executive (2002) and ASTM Standard G16 (2004).

The extreme value literature is extensive especially focusing on temporal extremes given an extensive historical data set of regular observations, for example Coles and Tawn (1994) consider a multivariate model using a limiting Poisson process while Coles and Tawn (1996) apply Bayesian methods using a limiting Poisson process to the extreme value problem. Spatio-temporal extremes are considered in Barão and Tawn (1999) in the context of a bivariate model which uses multivariate observations to make a single prediction based on asymptotic results. An extreme value analysis is carried out on corrosion data based on asymptotic results in Laycock *et al.* (1990) where measurements are of maximum pit depths. Environmental problems which are commonly temporally rich in data have motivated an extensive spatio-temporal literature. A dynamic linear model (DLM) framework is a common approach used in these circumstances to update uncertainties about general model parameters as observations are made by Shaddick and Wakefield (2002) and Stroud *et al.* (2001). Little *et al.* (2004) discuss the benefits of efficient inspection schemes and the potential reductions in inspection burden. Possible ways of reducing the computational load of the DLM include, for example, reducing dimensionality or constructing the update using summary variables. Systems that are strongly spatially correlated and hence are relatively smooth across space are best suited to this sort of technique (Winkle and Cressie, 1999).

In this paper, we describe the use of spatio-temporal DLMs for designing efficient inspection schemes for industrial systems which are corroding, where inspection data is aggregated and transformed, and design choice is analytically intractable. We proceed as follows. In Section 2, a motivating example is described. In Section 3, a spatio-temporal DLM which can be used to model continuous systems is considered. In Section 4, we discuss the application of Bayes linear methods to transformed data. In Section 5, we consider how to find optimal inspection schemes and the efficiency savings that can be made. As optimization of design choice is analytically intractable, we develop an efficient simulation methodology for the purpose. In Section 6, the application of the model to inspection planning is considered and its value is assessed. In Section 7, we discuss the potential benefits of the approach.

2 Motivating example

This paper concerns the benefits of careful modelling and analysis of spatio-temporal phenomena such as wall thickness. The work is motivated and illustrated by two cylindrical storage tanks. The first of these tanks stores caustic soda (sodium hydroxide solution), while the second tank stores crude acetone (2-propanone). Both caustic soda and acetone have relatively low corrosion properties, hence the tanks do not need frequent internal inspection and have a long expected life. The tanks were commissioned in the 1940s and are still in use at the current time. This work is not a case study of either of these tanks. However, we use the common physical aspects to motivate and illustrate the approach.

It is known that the lower part of these tanks will be more prone to internal corrosion than other parts of the tanks. In the example, only the lower part of the tank is modelled assuming all other parts of the tank will be less seriously corroded. There have been four inspections of each tank. Each inspection measures the wall thickness at a range of locations to assess vessel integrity at the time of the inspection. Each inspection is made using a different device and recorded in a different manner. The lack of detail in these historical data sets means making detailed use of them is very difficult.

Corrosion can be either internal or external. External corrosion is mainly due to rain and can be seen as small patches of external rust. This is found by regular external visual inspections and is then patched externally as necessary and painted over. The resulting wall may be thicker than might otherwise be expected. We do not have spatial details of the external repairs carried out. However all the repairs are small (less than 2 in. high) and located just above the floor. Internal corrosion will generally be more serious and require more substantial repairs – no repairs have been carried out following internal corrosion on either tank. The mean wall thickness for tank 3006 in the years 1969, 1980, 1992 and 2002, respectively, were 6.87, 5.94, 6.17 and 5.69 mm.

Particular problems with this data include: no record of the tank's design specification probably due to the long time scale involved; a lack of detailed information about the inspection methods used especially for the early inspections; considerable evolution of the inspection techniques used; lack of accurate spatial information for spatial data; irregularly spaced observations; and high observational noise.

The model we present can cope with all the problems associated with this set of data but, as would be expected, would give rise to a very wide credibility bound. Furthermore, using this data given the very limited background information and the extensive changes in inspection technology would be uninformative. Current expert beliefs are therefore used to motivate and construct a general form for the model. These problems emphasize the importance of having a flexible model which separates out the fine-scale description of the spatio-temporal development of the system from the details of how the system's integrity is checked. In this paper we optimize future inspection schemes based on current expert opinion about these and other similar tanks and current inspection technologies. The two tanks discussed in this section provide motivation for the illustrative example which considers the life cycle of a uniform tank.

3 Spatio-temporal DLMs

To predict the remnant life of a system, the remnant life of every part of the system as well as their covariance structure must be estimated using careful modelling and inspection. The inspection burden can be reduced by exploiting the spatial correlation of the system. If there is corrosion with no spatial correlation, 100% inspection may be the only option. Without a spatial model for the system it is necessary to make 100% inspection or accept that we have no knowledge about the locations where no inspection has been made. Similarly a temporal model is required if inferences are to be made about the future.

A Bayesian spatio-temporal model is a natural way to model many industrial systems with a continuous spatial element that develops through time, where there is extensive experience on the part of the operators but a relatively sparse amount of data. We follow the notation of the multivariate DLM developed by West and Harrison (1997).

Experts believe, if conditions remain constant, that current corrosion rate is an unbiased estimator of future corrosion rate, and that it is reasonable to assume that the system will deteriorate continuously while the tank is in use, so the corrosion is modelled with a locally linear trend. The system is modelled with two coupled equations: a system equation and a system slope equation.

$$\begin{aligned} \theta_{ts} &= \theta_{(t-1)s} + \beta_{ts} + \omega_{\theta ts} \quad && \omega_{\theta ts} \sim \langle 0, W_{\theta s} \rangle \\ \beta_{ts} &= \beta_{(t-1)s} + \omega_{\beta ts} && \omega_{\beta ts} \sim \langle 0, W_{\beta s} \rangle \end{aligned} \tag{3.1}$$

where θ_{ts} is the actual wall thickness and β_{ts} is the corrosion rate at time t and location s. $\omega_{\theta ts} \sim \langle 0, W_{\theta s} \rangle$ is used to denote a random vector with mean 0 and a variance matrix $W_{\theta s}$. In this example $W_{\theta s}$ and $W_{\beta s}$ are constant in time but dependence on t can be introduced if it is considered appropriate. It is further assumed that, given $W_{\theta s}$, $\omega_{\theta ts}$ is correlated over s for a given t but is independently and identically distributed (IID) for different values of t and similarly, given $W_{\beta s}$, $\omega_{\beta ts}$ is correlated over s, but IID over t. Given $W_{\theta s}$ and $W_{\beta s}$, $\omega_{\theta ts}$ and $\omega_{\beta ts}$ are mutually independent. $\omega_{\theta ts}$ 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. Note that $\text{var}(\theta_{ts} - \theta_{ts'})$ increases with time, meaning the surface becomes rougher.

It is also assumed that the prior beliefs about the initial system level and slope are in a similar form $\theta_{0s} \sim \langle 0, C_{\theta 0s} \rangle$ and $\beta_{0s} \sim \langle 0, C_{\beta 0s} \rangle$. It is assumed that there is no correlation between $\omega_{\theta ts}$, $\omega_{\beta ts}$, θ_{0s} and β_{0s} given $W_{\theta s}$, $W_{\beta s}$, $C_{\theta 0s}$ and $C_{\beta 0s}$, respectively.

Point observations of the system are modelled with an observation equation:

$$y_{ts} = \theta_{ts} + v_{ts} \quad v_{ts} \sim \langle 0, V_{ts} \rangle \tag{3.2}$$

where y_{ts} is the observed wall thickness value at time t and location s. V_{ts} may be time dependent (as would be required for the data described in Section 2) or fixed. It is further assumed, given V_{ts}, that v_{ts} is correlated over s for a given t but is IID for different values of t. v_{ts} represents observation and calibration errors which will not have any effect on the underlying wall thickness. For simplicity in this paper it is supposed that v_{ts} is independent and identically distributed noise but correlations over s for a given t can be introduced if required. It is also assumed, given V_{ts}, that v_{ts} is uncorrelated with all other quantities in the model. V_{ts}, $W_{\theta s}$ and $W_{\beta s}$ can be considered known, and then v_{ts}, $\omega_{\theta ts}$ and $\omega_{\beta ts}$ can be modelled using the Gaussian distribution, or the Bayes linear approach may be taken, considering v_{ts}, $\omega_{\theta s}$ and $\omega_{\beta s}$ as quantities for which prior beliefs are held which are specified by zero expectations and variance matrices V_{ts}, $W_{\theta s}$ and $W_{\beta s}$. For more information about the Bayes linear approach, see Goldstein (1999). Where observations are of the form of y_{ts} as mentioned earlier, updates of beliefs may be made in equivalent closed form by either approach.

There are many possible correlation structures for the system update variance matrix. In this example

$$\mathrm{cov}(\omega_{\kappa ts}, \omega_{\kappa ts'}) = W_{\kappa ss'} = \tau_\kappa \exp(-c_\kappa d_{(s,s')}) \quad \kappa = \theta, \beta \tag{3.3}$$

is used, where τ_κ is a measure of temporal strength, $c_\kappa > 0$ is a relative scaling parameter controlling spatial correlation and $d_{(s,s')}$ is a measure of distance between s and s'. Prior beliefs for θ and β may be based on any valid correlation structure. For simplicity the same correlation structure [Equation (3.3)] is used between the system levels ($W_{\theta s}$), the system slopes ($W_{\beta s}$), the prior levels ($C_{\theta 0s}$) and the prior slopes ($C_{\beta 0s}$) with parameters $\tau_{C\theta}$ and $\tau_{C\beta}$. Each of these four structures then has its own variance (defined by τ_θ, τ_β, $\tau_{C\theta}$ and $\tau_{C\beta}$, respectively). Possible ratios between these level and slope variances are discussed in Section 6.1. Were more detailed belief structures elicited for any of the four parts, they would be used instead. Using the model, means, variances and covariances can be calculated for any locations at the current time and in the future. The higher the correlation across the system the more information an observation will provide about the whole system.

3.1 Model parameters

The parameters used in the model are briefly described in this section. In subsequent sections the sensitivity of the model to changes in these parameters is analysed and discussed.

Within the example, only variance measures are considered. These are not affected by the mean wall thicknesses or the corrosion rate. In an application the initial wall thicknesses and corrosion rate can be combined with a minimum allowed wall thickness to give the expected life of the system.

The covariance between any two points (h_i, v_i) and (h_j, v_j) where h_i is the horizontal clockwise circumference distance from north and v_i is the vertical height and P is the total circumference in the tank may be described using

$$\begin{aligned}&\mathrm{cov}(\omega_{\kappa t(h_i, v_i)}, \omega_{\kappa t(h_j, v_j)})\\ &\quad = \tau_\kappa \exp\left(-c_\kappa \sqrt{(\min(|h_i - h_j|, \|h_i - h_j| - P|)^2 + a^2(v_i - v_j)^2)}\right)\end{aligned} \tag{3.4}$$

Different beliefs about horizontal and vertical correlation strengths may be reflected by rescaling the distance metric [parameter a in Equation (3.4)]. Following consultation with experts, for this example $a = 2\sqrt{2}$. The two key aspects of the problem, the spatial correlation and the temporal correlation, can be controlled by two parameters in Equation (3.4): spatial strength (controlled by c_κ) and update variance (controlled by τ_κ). For the purposes of this example, the prior level variance $\tau_{C\theta}$ can be fixed at 10 and the other parameters scaled accordingly using different values of τ_κ and $\tau_{C\beta}$.

3.2 Observation of minima

Simple DLM updating can be used where an observation is a true value with random noise. For more complex transforms of the underlying values, a more sophisticated analysis is required. In this section, we consider observations in the form of regional minima.

The distribution of the minimum of n independent $N(0,1)$ draws is skew. The severity of the skewness depends on the level of correlation between the data points. Analytical distributions and partial asymptotic results can be obtained for minima of uncorrelated systems but results are less tractable for correlated distributions. In the example, a sample that covers only a small fraction of the system is used as a basis for establishing the entire system's integrity using the spatial correlation.

The true wall thickness is continuous and this continuum is modelled with a finite set of correlated points. Fewer points result in a lower model accuracy but faster computations. Similarly, higher correlations also reduce the number of points required to gain a given credibility level in the summary observations. The accuracy of the approximation made when modelling a continuous system with a discrete grid will increase as a higher density of points is modelled. High correlation and low observation noise will also help increase model accuracy. For any set of parameter values the sensitivity to changes in the density of points should be investigated.

The lower part of the tank is modelled with n_h equally spaced points horizontally around the tank, n_v equally spaced vertically orientated points in the lower part of the tank and n_t equally spaced time points. For this example, $n_h = 640$, $n_v = 5$ and $n_t = 8$. An observed regional minimum is defined as the minimum of a 5 (horizontal) by 5 (vertical) subsection of the grid at a particular time which includes point by point IID observational noise, giving a maximum of 640 possible different regional minima for any inspection. The observed regional minimum at time t and centred at location $s = (h, 3)$ may be written as $\mu_{yts} = \min_{i,v} y_{t(i,v)}$ where $i = h-2, h-1, \ldots, h+2$ and $v = 1, 2, \ldots, 5$ where y_{ts} is defined in Equation (3.2). Similarly, the underlying regional minimum at time t and centred at location s may be written as $\mu_{\theta ts} = \min_{i,v} \theta_{t(i,v)}$. The effect on the estimated regional minimum value of increasing the density of points in each region should be carefully considered. An acceptable approximation of the continuous surface should not be sensitive to small changes in this density.

4 Bayes linear forecasts

This paper develops efficient inspection schemes for reducing uncertainty about system corrosion with minimum inspection cost. To identify efficient inspection schemes many candidate schemes must be compared, where both the number of times at which to inspect, and also the number of locations to inspect at each time point may be varied. A full Bayes preposterior evaluation of the reduction in uncertainty associated with any one inspection scheme based on data in the form of regional minima would be enormously computationally intensive. Where a range of design choices are to be considered, it is important to use an approach which can be evaluated reasonably quickly for each possible design, and the Bayes linear approach is well suited to such problems.

Analysing design choice problems based on transformed data is challenging because transformed observations are often associated with distributions which are analytically intractable. The correlations between observed (transformed) values and transforms of the system values in which interest is focused can be calculated using simulation. Beliefs about these system-based transformations are then updated using a Bayes linear framework. Critical issues, such as identifying when and where future inspections should be made, and questions about the confidence in the system's integrity in the future can be addressed within this framework.

4.1 Forecasting

In this section, the forecasting of transformed outputs (e.g., regional minima) at time $t + k$ when transformed observation values are observed at time t is discussed. Forecasts are made directly, from the data to the quantities over which inferences are required (for example, the global minima of the system values, at certain future time points). A suitable update can then be used to adjust the forecast mean and variance at time $t + k$ given the observations up to time t. Using a fully Bayesian update for minima data would be very computationally demanding. The Bayes linear approach uses simple updating equations to calculate the adjusted expectation

$$E_{D_t}(\phi) = E(\phi) + \mathrm{cov}(\phi, D_t)\mathrm{var}(D_t)^{\dagger}(D_t - E(D_t)) \tag{4.1}$$

and the adjusted variance

$$\mathrm{var}_{D_t}(\phi) = \mathrm{var}(\phi) - \mathrm{cov}(\phi, D_t)\mathrm{var}(D_t)^{\dagger}\mathrm{cov}(D_t, \phi) \tag{4.2}$$

where $\dagger$ is a generalized inverse and ϕ is a vector whose values it is intended to learn about via D_t the vector of all observed data up to time t. In this application D_t is a vector of terms $\mu_{yt_is_i}$ while ϕ is a vector of terms $\mu_{\theta t_js_j}$. In particular $\mathrm{var}_{D_t}(\phi)$ depends only on the observation locations and not on the observed values themselves and so can be used as the basis of a tractable design methodology. Using the Bayes linear approach, the forecast values for future observations and any underlying system parameters (for example true regional minima) are adjusted directly following an observation. The covariance between all the observed data values and all other system quantities of interest at all times must be calculated (Section 4.2). There is no requirement to have the same method of data collection or the same measurement variance at each time point.

4.2 Simulation

Simulation may be used to construct realizations of the system consistent with given prior beliefs. Simulated data can be used to estimate expectations, variances and covariances where it would be difficult or impossible to calculate these by analytical means. For example, if the k step ahead covariance between two regional minima is required full data sets for both regions can be simulated repeatedly and hence the covariance between the minima of the two regions estimated. Forecasts require correlations of the underlying minima at the time and location of the forecast with the observed minima for each observation that has been made. Hence, to do general design based calculations, covariances are required between all observed minima and underlying minima across all times and regions of interest.

If the problem is too large for direct calculation of the covariance matrix from simulated data the symmetry of the problem may be exploited. The covariance between any pair of minima which are the same spatial distance apart at fixed times will be the same. That is, $f(d, t, t') = \text{cov}(\mu_{\theta ts}, \mu_{\theta t's'})$ where d is the distance between s and s', and similarly for $\text{cov}(\mu_{yts}, \mu_{yt's'})$ and $\text{cov}(\mu_{yts}, \mu_{\theta t's'})$. A single calculation of this covariance can then be used for all observed regional minima or underlying regional minima between two regions at time t and t' that are d spatial units apart. This can then be repeated for all d, t and $t' \geq t$ and hence the full covariance structure found. While this is fast and efficient, for large problems (e.g., if all 640 possible minima must be considered) due to simulation error the resulting covariance matrix may not be positive definite, and it would be necessary to reconstruct the variance matrix for example with negative eigenvalues removed.

For the level of detail described in this problem, a full simulation of data across the entire observed and underlying systems is practically possible but slow. For this example only 64 regional minima are considered evenly spaced horizontally around the tank. This means only half of the horizontal points on the grid are considered in sets of five, with five points observed and five points unobserved. These 64 regional minima are assumed to have horizontal midpoints at $1, 11, \ldots, 631$ and are, respectively, indexed by $\psi_1, \psi_2, \ldots, \psi_{64}$ which form a set Ψ. Henceforth only the observed regional minima $\mu_{yt\psi_i}$ and underlying regional minima $\mu_{\theta t\psi_j}$ where $\psi_i, \psi_j \in \Psi$ are considered. This smaller number of minima also reduces the computational load of subsequent searches to find the optimal set of inspection locations (Section 5). These searches are very demanding when there is a large number of possible combinations to consider. Once design choices have been made, these intermediate regions should be modelled to check for unacceptably high posterior variance.

For this example a set of initial wall thicknesses are simulated in accordance with prior beliefs θ_{0s}. Level and slope updates ($\omega_{\theta ts}$ and $\omega_{\beta ts}$) are then simulated for times 1–8. The prior and update simulations are then combined to get true wall thicknesses θ_{ts} at the 1600 locations modelled over the 8 times. The 64 underlying regional minimum wall thicknesses ($\mu_{\theta t\psi_i}$) and the 64 observed regional minimum wall thicknesses ($\mu_{yt\psi_j}$) are calculated for each time (the latter by adding the IID observation noise v_{ts} to the 1600 underlying wall thicknesses). Finally the correlation structure between the underlying regional minima and observed regional minima wall thicknesses across times 1–8 are calculated over 40 000 simulations of this data. Having calculated the covariance between all observed regional minima and underlying regional minima of interest across all times, the effect of making any particular set of inspections on the uncertainties across the entire system may be assessed.

5 Finding the optimal scheme

The inspection design Δ is defined by i) the vector m_t of times at which inspections are made and ii) the locations at which inspections are made at time t, $\psi_{t_1}, \ldots, \psi_{t_{m_{nt}}}$, where $t_1, \ldots, t_{m_{nt}}$ are m_{nt} choices from the index set $[1, \ldots, 64]$. In this example, m_{nt} is restricted to be the same value m_n for each time at which inspections are made. If two inspections are made over four different times, with 64 possible locations for each inspection, the total number of different inspection schemes is approximately $(64^2)^4$. Extending this to 8 time points gives approximately 7.9×10^{28} distinct inspection schemes to be considered. Efficient ways of reducing the number of possible alternatives must be considered. The simplifying restriction is made that, if m_n regional minima inspections are made at time t, then the m_n inspections are equally spaced around the tank. This gives rotational symmetry. In addition to being an intuitively natural simplification, the symmetry drastically reduces the number of possible locations that need to be considered for future inspections. A more complex inspection scheme might try and learn about corrosion correlation across different locations but this is a different modelling problem.

It is necessary to maintain a certain level of confidence in the system's integrity at all times. At any given time before observations are made, the regional minimum wall thicknesses will have a constant expected value across space. A sensible example criterion for selecting inspection points is now considered although many other criteria are possible and would be evaluated in a similar fashion. The criterion considered is the average of the maximum adjusted variances of the underlying regional minimum wall thicknesses from all the locations modelled, Ψ, over all times up to forecast time T for a given inspection design Δ which is denoted $F(\Delta)$, where $F(\Delta) = (1/T)\sum_{t=1}^{T} \max_{\psi_i \in \Psi} \text{var}_{D_{t-1}}(\mu_{\theta t \psi_i})$, where $\mu_{\theta t \psi_i}$ is the underlying regional minima at time t and location ψ_i, and D_t is all observed data up to time t. $\text{var}_{D_{t-1}}(\mu_{\theta t \psi_i})$ is found using Equation (4.2).

The inspection design Δ will be restricted by the times when inspections are allowed (m_t) and the number of locations inspected (m_n) when an inspection is made. For any choice of m_n and m_t, the optimal inspection scheme $\Delta^*(m_n, m_t)$ will then minimize the criterion $F(\Delta)$ over all designs Δ with the given values m_n, m_t and rotational symmetry at each time point.

5.1 Possible search strategies

Having restricted attention to equally spaced inspections, for small problems a full search can be made. In larger problems, further load reducing strategies are required. The symmetry of the prior beliefs may be exploited to reduce the number of inspections that must be considered. A 'tentacle' search from likely solutions may also be used. Given an initial solution, small changes to the locations inspected at each time are considered, repeating the search with the improved solution until convergence occurs. An alternative is initially only to assess some of the possible inspection location combinations before focusing in more detail on the locations with the best preliminary results. A 'tentacle search' can then be made around any possible minima. A further alternative is instead of simultaneously choosing inspection sites across all time points simultaneously, the temporal horizon of the model may be reduced. All possible inspections between times 2 and $k+1$ may be considered. The best of these candidate schemes is then used to fix the inspection location at time 2. All inspection locations between times 3 and $k+2$ are then considered and the best of these candidate schemes used to fix the inspection location at time 3. This may then be repeated until all times have been optimized over. The larger the k the better this scheme will be, but also the greater the computational load. An increase in time points will only increase computational load linearly. These methods can be combined to produce an efficient search for a good inspection scheme. In the following section a full search is used where reasonably possible, and otherwise a 2 step ahead search is applied to a 50% grid followed by a tentacle search until a local minimum is found.

6 Application to inspection planning

In this section the effect of the covariance and the prior belief specification on the output of the model are considered. A system is developed to provide decision support when assessing how often and at how many locations inspections should be made. Information is presented in a graphical format allowing easy comparison between possible schemes and an assessment of model sensitivity.

6.1 Correlation between minima

Maximum temporal correlation is achieved when there is little update uncertainty. Under these circumstances the prior slope variance ($C_{\beta 0}$) will dominate forecast variance for sufficiently large t. Within the framework considered, prior slope variance ($C_{\beta 0}$) is a fixed proportion of prior level variance ($C_{\theta 0}$). The prior level variance is fixed at 10 (Section 3.1) and hence this ratio is critical if high temporal correlations are to be considered.

For any $W_{\theta s}$, fixing the ratio of $W_{\theta s}$ to $W_{\beta s}$ at $10:1$, the temporal correlation between an observed regional minimum at time 1 ($\mu_{y1\psi_i}$) and the true regional minimum over the same area at time 8 ($\mu_{\theta 8\psi_i}$) has a maximum correlation strength of about 0.4. Increasing the ratio to $100:1$ the maximum correlation reaches about 0.8 while at ratio of $1000:1$ the correlation can reach almost 1. Small temporal correlations will still occur when update uncertainty is large. Henceforth a ratio of $1000:1$ is assumed to allow a full range of correlations.

With a temporal strength (τ_θ) of 1 and a spatial strength (c_κ) of 0.001 these beliefs result in a spatial correlation at time 1 between a regional minimum observation of one region (ψ_i) and the true regional minimum of the region on the opposite side of the tank ($\psi_{(i+32)}$) of about 0.7 and a temporal correlation between an observed regional minimum at time 1 (ψ_i) and the true regional minima at time 8 (ψ_i) also of about

0.7. Different pairs of spatial and temporal parameters are now taken around these values to explore the sensitivity of the inspection plan to different strengths of spatial and temporal correlation. As shown in Table 1, temporal parameter values of 5 and 0.3 and spatial parameter values of 0.0025 and 0.0003 give temporal and spatial correlations (as defined earlier) of approximately 0.5 and 0.9. These correlations (ρ) were chosen to approximate a three by three grid of ρ^2 values based on approximately 0.25, 0.5 and 0.75 which are now considered for both temporal and spatial components. The first value of each pair in Table 1 is the temporal correlation while the second value in each pair is the spatial correlation both described earlier. IID observation noise has a standard deviation of 1.0.

We have found low levels of correlation in the observation noise have little effect on the model. Since observational noise has little effect on the correlations, any forecasts made will also be reasonably robust to changes in the observational noise. Observation noise with a standard deviation of 1.0 is relatively large compared with the standard deviation of the general temporal component (0.3, 1 or 5) but less so compared with the prior level uncertainty (10).

6.2 Comparison between inspection schemes

Suppose that for a given inspection scheme there are m_n equally spaced observed regional minima ($\mu_{yt_i\psi_j}$) during each inspection. Different numbers of inspections are considered; $m_n = 1, 2, 4, 8$ and 16. Having chosen the first inspection, a range of possible locations for subsequent inspections may then be considered. For each inspection regime, the adjusted variance of each underlying regional minima may be assessed at each time point. If at any time point, sufficient corrosion is discovered to raise doubts about the future integrity of the vessel then it will be necessary to repair the vessel, or, at least increase the frequency of inspection around the region causing concern. Five different inspection schemes are considered (Table 2) with m_n observations made at each time.

The proportion of the initial uncertainty remaining (given prior knowledge D_0) following observation of the data may be calculated. Taking this proportion away from 1 gives the proportion of variance removed by the inspection scheme. This may be written as

$$P(m_n, m_t) = 1 - \frac{F(\Delta^*(m_n, m_t))}{\max_{\psi_i \in \Psi} \mathrm{var}_{D_0}(\mu_{\theta T \psi_i})}$$

where $T - 1$ is the time of the last inspection.

Table 1 Temporal and spatial correlation of minima

Temporal (τ_θ) Spatial (c_κ)	5	1	0.3
0.0025	0.44, 0.45	0.72, 0.44	0.87, 0.44
0.001	0.43, 0.71	0.73, 0.71	0.88, 0.72
0.0003	0.44, 0.89	0.73, 0.89	0.87, 0.89

Table 2 Different inspection schemes

Scheme (number of times inspected)	Times inspected (m_t)
a(1)	1
b(2)	1, 7
c(3)	1, 4, 7
d(4)	1, 3, 5, 7
e(7)	1, 2, 3, 4, 5, 6, 7

Figure 1 shows for each m_n, set of inspection times (m_t) and set of values of τ_θ and c_κ the value of $P(m_n, m_t)$ for design $\Delta^*(m_n, m_t)$. Near white indicates an informative inspection procedure while near black indicates an uninformative inspection procedure. The middle subgrid corresponds to spatial and temporal parameters of 1 and 0.001, respectively. Each square in this grid then corresponds to a given number of inspections at each time point (on the horizontal axis) and inspections made at times indicated by the vertical axis. Locations are chosen as in Section 5 to optimize $F(\Delta^*(m_n,m_t))$.

The top right subgrid has the maximum spatial and temporal correlation which is reflected by the high proportion of the available information that is gained by any inspection. The increase in information following increased temporal or spatial inspections is shown by the larger values in the top right of each subgraph for all parameter values. Using these results it is possible to compare the benefit of extra spatial inspections to the benefit of extra temporal inspections. An inspection cost function may be used to make a comparison between costs and gain in accuracy for different inspection schemes. In its simplest form this function will include a cost for each regional minimum observed and a fixed cost for making one or more inspections at

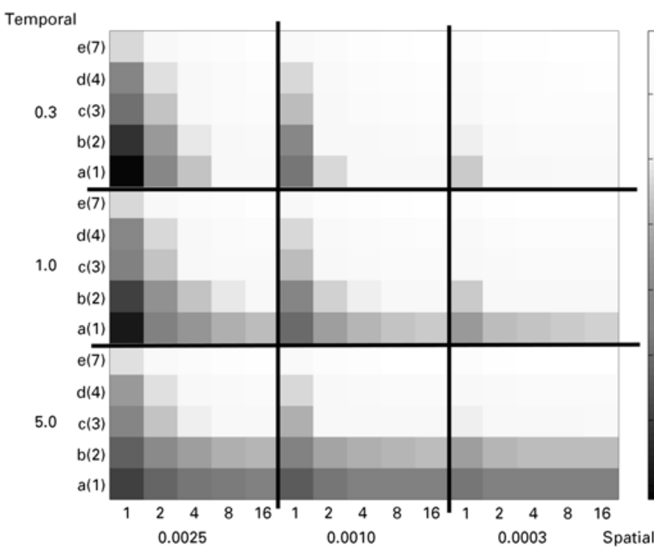

Figure 1 Matrix of local simulations

a given time. The scheme that gains maximum information for a fixed cost or a given amount of information at minimum cost can then be found, and the sensitivity of the design choice to the correlations in the model explored.

6.3 How much improvement

Inspection locations have been selected based upon reducing system variances. These spatial schemes may be compared to the simpler nonspatial schemes with the same number of inspections and fixed inspection locations through time (required for nonspatial models) to see how much additional information is gained. The benefits for the nine pairs of parameters (introduced in Table 1) when an inspection is made at every time interval with a fixed number of points (1, 2, 4, 8 or 16) relative to the same number of points but with fixed locations can be seen in Figure 2. The large values to the left of the plot indicate significant information gains potentially reducing variance up to 3.5-fold when a suitable inspection scheme is used. The trend down to the right shows that as more points are inspected having variable inspection locations become less important. Lower temporal effects generally make the inspection locations more important while, conversely, greater spatial strength generally reduces the importance of the inspection locations. Also note that, in this example, making repeat inspections at $2n$ locations is typically as effective at reducing variance as making n inspections at selected locations (not shown in Figure 2).

6.4 Prediction of global minimum

The life of the system is essentially defined by the global minimum. The variance of the global minimum may be adjusted for a particular sampling scheme based upon ranked

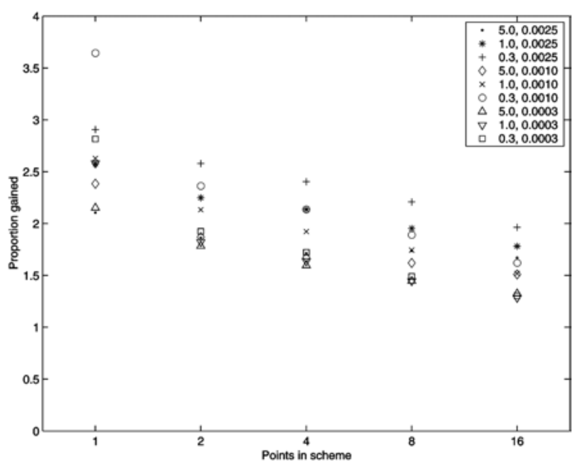

Figure 2 Gain from spatially informative sampling

regional observations as discussed in what follows. The expected global minimum can be adjusted likewise.

For a particular inspection scheme, the prior covariances between 'the observed regional minima after they have been ranked' and the 'true global minimum at each time point' are calculated. These covariances are calculated directly from the data previously simulated to find the covariance structure between the regional minima. The covariance structure is then adjusted using Equation (4.2) to evaluate the effect of a set of regional minima observations on beliefs about the global minimum.

A particular scheme is considered with temporal covariance parameter $\tau_\theta = 1$, spatial covariance parameter $c_\kappa = 0.001$ and $m_n = 4$ equally spaced inspections made at times 1, 4 and 7. Four equally spaced regional minima are thus observed at three different times. There are 16 candidate inspection locations on each quadrant of the tank and within each quadrant inspections are made at candidate locations 1, 7 and 14 (at times 1, 4 and 7, respectively). These candidate locations were found to reduce the one step ahead global system variance following adjustment by the ranked observations the most over all possible schemes based on a small subset of the data (1000 cases). Note that for this example, using the methods previously described, 89% of the initial uncertainty over the regional minima can be resolved. The covariance is now found after ranking by observation size rather than by finding the covariance directly between these 12 inspections and the global system minimum at each time. In this example on average over the eight times approximately 6% of the prior variance about the one step ahead global system minimum remains following adjustment by the ranked observations. Using the same data but not ranking leaves approximately 50% more of the initial uncertainty about the global system minimum.

7 Discussion

We have discussed efficient inspection for maintenance of complex industrial systems subject to degradation effects, such as corrosion. With appropriate statistical modelling, the potential for better use of inspection resources to improve the quality of inferences has been demonstrated. A Bayesian spatio-temporal DLM has been developed suitable for a wide range of problems, such as wall thickness monitoring, where the inspection method used collects transformed data. Historically the results of existing inspections have often been found to be badly recorded. It has been shown that detailed recording of inspection locations coupled with a suitable spatio-temporal model may substantially reduce uncertainty following an inspection.

The developments made in this paper can be used in a decision support role during the inspection of complex industrial systems. Our approach allows efficient inclusion of transformed observations which are difficult to use directly to update beliefs in design calculations for the standard DLM setup. This allows us to construct efficient designs under a variety of different sampling schemes.

Having identified a good inspection design for a given uncertainty specification, the effect of changes in prior uncertainty on the preferred inspection scheme should be assessed. If feasible changes in the prior beliefs governing the model do not significantly change the preferred inspection locations for a given inspection budget, then the design is robust to changes in parameters. A sensitive model will require consideration of the additional risk associated with plausible changes in parameter values and may suggest that a more extensive inspection is required.

In order to devise an optimal sampling scheme, the probabilistic information that we display must be combined with a careful specification of the utilities involved in comparing sampling cost with the consequences of system failure. This paper provides the relevant probabilistic information to allow the manager to make an informed choice between competing inspection schemes. More sophisticated design choices could be made aiming at learning both about the correlation parameters and also the likely future behaviour of the system, but, analysis of such designs is beyond the scope of the paper.

Having identified a suitable design, the density of points used to evaluate regional and global minima should, if computationally possible, be increased to check for sensitivity. If using a higher density of points has a significant effect on the expected outcomes then the density of points used for the calculations may have been too low. This will be particularly important when spatial correlations are low. Finally careful model checking, after each inspection, is essential. An efficient collection of one step ahead forecast diagnostics for such purposes is described in Little *et al.* (2004).

References

ASTM Standard G16 (2004) *Standard guide for applying statistics to analysis of corrosion data*. PA, USA: ASTM International.

Barão MI and Tawn JA (1999) Extremal analysis of short series with outliers: sea levels and athletics records. *Applied Statistics* **48**, 469–87.

Coles SG and Tawn JA (1994) Statistical methods for multivariate extremes: an application to structural design. *Applied Statistics* **43**, 1–48.

Coles SG and Tawn JA (1996) A Bayesian analysis of extreme rainfall data. *Applied Statistics* **45**, 463–78.

Goldstein M (1999) Bayes linear analysis. In Kotz S, Read CB and Banks DL, editors, *Encylopaedia of statistical sciences*. Update Volume 3. New York: Wiley, 29–34.

Laycock P, Cottis R and Scarf P (1990) Extrapolation of extreme pit depths in space and time. *Journal of Electrochemical Society* **137**, 64–69.

Little JD, Goldstein M, Jonathan P and den Heijer C (2004) Spatio-temporal modelling of corrosion in an industrial furnace. *Applied Stochastic Models in Business and Industry* **20**, 219–38.

Shaddick G and Wakefield J (2002) Modelling daily multivariate pollutant data at multiple sites. *Applied Statistics* **51**, 351–72.

Stroud JR, Muller P and Sansó B (2001) Dynamic model for spatiotemporal data. *Journal of the Royal Statistical Society* **63**, 673–89.

UK Health and Safety Executive (2002) *Guidelines for use of statistics for analysis of sample inspection of corrosion*. Norwich. [http: //www.hse.gov.uk/research/rrpdf/rr016.pdf (last accessed 21 October 2004)]

West M and Harrison J (1997) *Bayesian forecasting and dynamic models*, 2nd edition. New York: Springer-Verlag.

Wikle CK and Cressie N (1999) A dimension-reduced approach to space-time Kalman filtering. *Biometrica* **86**, 815–29.