JOURNAL OF MATERIALS SCIENCE 31 (1996) 5127–5134

Spatial statistics of particles and corrosion pits in 2024-T3 aluminium alloy

N. R. CAWLEY, D. G. HARLOW
Mechanical Engineering and Mechanics, Lehigh University, 19 Memorial Drive West, Bethlehem, PA 18015, USA

Corrosion pits in aluminium alloys nucleate and grow in the proximity of constituent particles. Descriptions of the spatial statistics of the particles and pits are critical for stochastic modelling of corrosion. Statistical models, which are generalizations of a hard-core sequential inhibition process, are proposed. The models include the effects of randomness in the number, the locations, and the sizes of the features. The applicability of the models are based on experimental data taken from 2024-T3 aluminium alloy specimens subjected to a 0.5 M NaCl environment. The observed constituent particles tend to be clustered, whereas the subsequent corrosion pits tend towards regularity as measured by nearest neighbour spacing.

1. Introduction

Pitting corrosion in aluminium alloys is recognized as one of the significant degradation mechanisms that impacts the reliability, durability, and integrity of both military and commercial aircraft. Pitting corrosion typically is the precursor to rather complex damage processes such as corrosion fatigue crack initiation and growth; however, it is a complex process in its own right. Corrosion pits nucleate and grow where constituent particles can serve as local galvanic cells. The nucleation and growth processes of pits are stochastic processes which depend upon many intricate material properties, the applied loading, and the environmental conditions. Thus, the random evolution of the number of pits and their sizes is an extremely complicated process. This random process can be considered as a special case of the general class of stochastic processes known as nearest neighbour processes. Nearest neighbour processes have been studied extensively in the statistics literature; see the books by Ripley [1], Diggle [2], Stoyan *et al.* [3], and Cressie [4], for example.

The conventional wisdom is that the materials which are weaker or have shorter times to failure are those for which excessive clustering of pits or very large pits occur. The purpose of this work was to adapt modelling techniques from nearest neighbour methods in spatial statistics to gain insights into the random locations of constituent particles and subsequent corrosion pits in aluminium alloys. Thus, the goal was to produce a methodology for statistically analysing the random position of particles and pits which can be incorporated into models for corrosion fatigue crack initiation and growth.

The material considered was 2024-T3 aluminium alloy. Each specimen was polished on a surface in the rolling direction, and its constituent particles were observed prior to exposure to the environment. The specimen then was placed in a 0.5 M NaCl solution for some prescribed time, after which the corrosion pits were observed and measured by optical microscopy. Five different specimens were considered: a typical polished specimen prior to any exposure to the environment, and four specimens exposed to the environment for 10, 24, 42, and 72 h, all at a temperature of 40 °C.

Typically polished sections of the 2024-T3 alloy, used in this work, were observed to contain approximately 3000 particles, with an area of at least 1 μm^2, per mm^{-2}. The particles are either anodic relative to the aluminium matrix, in which case they dissolved, or cathodic relative to the matrix, in which case the matrix dissolved. There were about three times more anodic particles than cathodic particles for the observations herein. Because the primary purpose of this work was to study the spatial statistics of the constituent particles and the subsequent corrosion pits, further material characterization was not required. A detailed discussion of the properties of this alloy is given elsewhere [5].

2. Spatial point processes

In the last two decades, there has been a flurry of activity in theoretical and applied statistics for spatial point processes [1–4]. A spatial point process is any stochastic process that generates a countable set of events, x_i, in a subset of a two- or three-dimensional space. The characterization of the locations of the constituent particles and subsequent corrosion pits in aluminium alloys is well suited as an application of spatial point processes. The first spatial point process to be considered is the complete spatial randomness (CSR) model. The CSR model implies that the pattern

for the centroids of the particles or pits are uniformly distributed over the target area. The basic assumption for this model is that each feature may be considered to be a single point, i.e. the area of the feature may be neglected. If the features are sufficiently disperse, or if their areas are sufficiently small, then this assumption may be reasonable; however, features from most applications rarely satisfy these conditions. The dominant reason for considering the CSR model is that the underlying stochastic process is a spatially homogeneous Poisson process which uniquely serves as a standard of comparison for all other spatial processes.

Fig. 1 contains typical micrographs of a polished specimen and a specimen after exposure to the corrosive environment. Even a casual observation of Fig. 1 is sufficient to reject the CSR model assumption that the feature may be considered simply as points consisting of no area. Thus, the CSR model is not appropriate. The next assumption typically included in modelling of spatial features is that each feature is shaped regularly, and the shape can be estimated by the observed cross-sectional area. Again, neither the particles nor the pits in Fig. 1 are shaped regularly. In fact, there are a variety of different shapes. Ultimately more general shapes should be considered in the model; however, the initial analysis will assume regularity.

During the evolution of an individual pit, it may continue to grow, or it may coalesce with another pit. Under any given conditions when an optical image is taken, it is impossible to determine the stage of the evolution process for the pits. If Fig. 1 is considered carefully, the features appear to be distinct, for the most part. Thus, a hard-core model may be appropriate to model the pit locations. A hard core model [1–4] is one in which the point pattern generated from the underlying stochastic process consists of non-intersecting areas. There are several types of hard-core models, but the form to be considered here is a generalization of the simple sequential inhibition (SSI) model [6]. Even though the SSI model is one of the more tractable models, it is still sufficiently complex that explicit statistical modelling is typically intractable. One way in which to approach the problem is to employ implicit modelling. An implicit model is one in which the probability distribution functions are intractable. A common method of statistical inference for an implicit model is through the use of Monte Carlo simulations [7]. In general, the implicit statistical model, described in enough detail, is simulated repeatedly, and the results are matched to experimental data. If the data conform to some prescribed statistical test, generated from the simulation, then the model is assumed to be representative of the underlying physical process. Obviously, the implicit statistical model should not be unduly complex. The more simple model which approximates the data is to be preferred. In the remainder of this paper, the general properties of spatial point processes needed in the modelling will be given, generalizations of the SSI model will be applied to the data, and a discussion with conclusions will follow.

3. Analysis of sampled patterns

The spatial point processes for constituent particles and corrosion pits are assumed to be stationary, invariant under arbitrary translations, and isotropic, invariant under arbitrary rotations. The practical implication of stationarity is that replication within a single set of data is possible in disjoint subsections of the target. Isotropy simply means that vector differences between events can be represented by scalars. Both of these assumptions are reasonable because the environment is uniform over the entire specimen.

The most commonly used properties to describe a stationary and isotropic spatial point process are the first-order and second-order properties. The first-

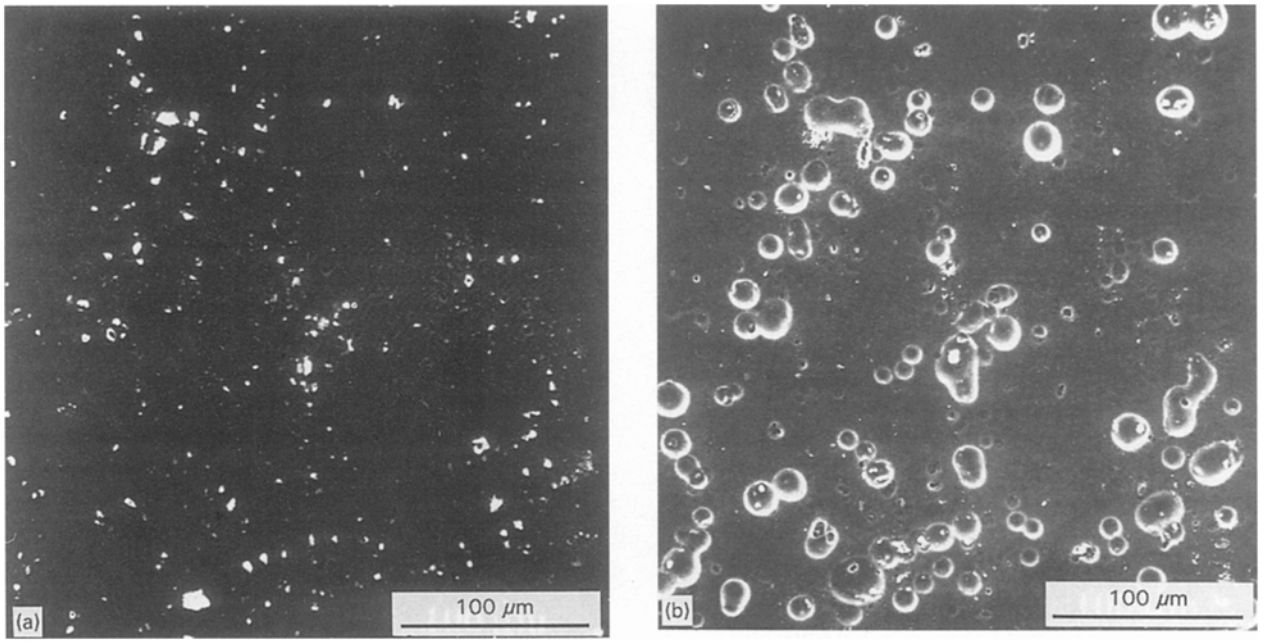

Figure 1 Typical scanning electron micrographs of 2024-T3 aluminium alloy (a) before and (b) after exposure to 0.5 M NaCl solution for 72 h at 40 °C.

order properties are given by the intensity, λ, which is defined as

$$\lambda = E(\text{number of events, i.e. features, per unit area}) \quad (1)$$

where $E[\ ]$ is the expectation operator. The second-order properties are given by the second reduced moment function, $K(t)$, one definition of which is

$$K(t) = E(\text{number of additional events within distance } t \text{ of an arbitrary event})/\lambda \quad (2)$$

Because λ and $K(t)$ are both defined as expectations, they do not completely characterize the underlying random process. Nevertheless, they do restrict the number of models which can serve to represent the data. The motivations for considering λ and $K(t)$ are two-fold. First, both λ and $K(t)$ can be estimated rather easily from data, and the statistic is superior to others used for point processes. Second, they are applicable for an extremely wide range of point processes.

Ripley [8] introduced an estimator for $K(t)$ which is effective for a wide range of applications. Let n be the number of observed events in a region A, and let $|A|$ be the area of the field of observation. Then the estimator for $K(t)$ is given as follows

$$\hat{K}(t) = n^{-2}|A| \sum_{i \neq j} \sum w_{ij}^{-1} I_t(u_{ij}) \quad (3)$$

where $I_t(u)$ is the indicator function defined by

$$I_t(u) = \begin{cases} 1, & \text{if } u \leqslant t \\ 0, & \text{if } u > t \end{cases} \quad (4)$$

The term u_{ij} is the distance between events i and j, and w_{ij} is the proportion of the circumference of the circle centred at event i with radius u_{ij} which is contained within A. Equation 3 is an approximately unbiased estimator for sufficiently small t because $n/|A|$ is a slightly biased estimator for λ. The restriction on t is necessary because $w_{ij}^{-1} \to \infty$ as t increases. In most applications, it is the local interactions between events that are of interest so that the restriction of small t, in fact, poses no problem. Fortunately, an explicit formula exists for w_{ij} when the field of observation is rectangular. Diggle [2] gives further details including the explicit equation required for w_{ij}.

Because second-order properties do not uniquely define a process, another description for the point process is in order. There are several candidates, but one of the more natural is the empirical distribution function (edf) $G(y)$ defined by

$$G(y) = Pr(\text{distance from an arbitrary event to its nearest neighbour is at most } y) \quad (5)$$

Ripley [9] introduced the following unbiased estimator for $G(y)$. For $1 \leqslant i \leqslant n$, let y_i be the distance from each event to its nearest neighbour in A, and let d_i be the shortest distance from each event to the boundary of A. Then the estimator for $G(y)$ is given by

$$\hat{G}(y) = \frac{\#(y_i \leqslant y \text{ and } d_i > y)}{\#(d_i > y)} \quad (6)$$

where $\#(\ )$ is the counting function which tallies the number of events in the specified set.

The primary test for the spatial statistics models for the constituent particles and the corrosion pits will involve a comparison of the estimate of $K(t)$ for the data with estimates of $K(t)$ from simulated realizations of the correspondingly proposed model. Secondarily, another test will be the compatibility of estimates of $G(y)$ from both the data and the simulated realizations. These tests are applied below in a series of progressively more complicated attempts at modelling the spatial statistics of the features.

4. Modelling the spatial statistics of constituent particles and corrosion pits

A few comments about the physical parameters which influence the nucleation and growth of corrosion pits are in order. Given a specimen, the statistics from a spatial point process can only model the specific pattern, and as such, the physical parameters are implicit in the modelling. However, it is manifest that other quantities have a profound effect on the kinetics of pit growth. In order statistically to correlate these types of key parameters with the observed spatial point process, several critical experiments should be conducted. The purpose of this work was to focus on modelling the observed point process, and consequently, modelling the entire evolution process is deferred.

In order to gain an appreciation for the spatial locations of particles and pits, Fig. 2a shows a schematic representation of the particles observed on a polished specimen, and Fig. 2b is similar for the corrosion pits on another specimen, which had been exposed to 0.5 M NaCl solution for 42 h at 40 °C. The figures were constructed by placing a circular disc, centred at the centroid of the particle or pit, with area equal to the measured area. Although neither the particles nor the pits are actually circular in shape, Fig. 2 is intended to illustrate the randomness in the number, size, and location of the features observed on typical 2024-T3 specimens. In both representations there are a few features which appear to intersect. This is an artefact of the assumption of circular features rather than reality. It is clear from Fig. 2 that constituent particles and corrosion pits exhibit quite different spatial patterns; however, notice that it is nearly impossible to describe adequately the patterns of the features from observations alone. Certainly, there is no apparent spatial structure. Thus, the need for statistical modelling is manifest.

The centroid data were used to calculate the value of $\hat{K}(t)$ from Equation 3 for the spatial point process, and the results for all five specimens are given in Fig. 3. For the CSR model, $K(t)$ can be found explicitly, and it is given by

$$K(t) = \pi t^2 \quad (7)$$

Thus, the vertical scale has been normalized so that the diagonal line is characteristic of the CSR model. Thus, the normalized reduced second moment measure is defined as $[K(t)/\pi]^{1/2}$. The four data sets for the corrosion pits all exhibit similar behaviour in

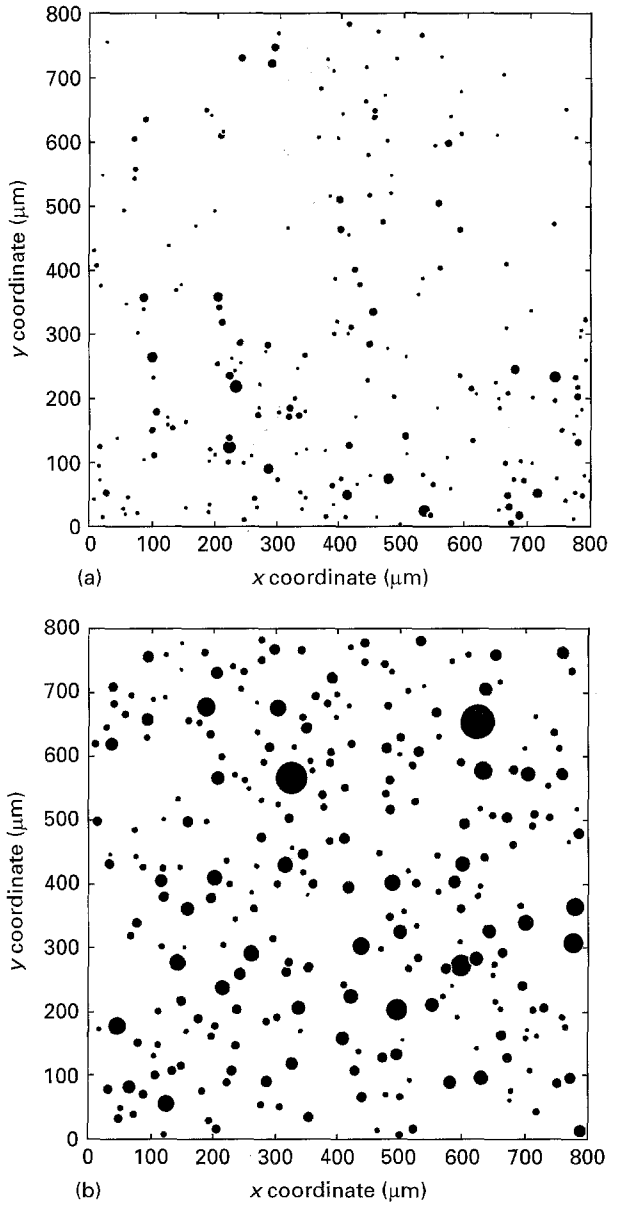

Figure 2 Schematic representations of 2024-T3 aluminium alloy specimens showing (a) constituent particles on a polished section and (b) corrosion pits formed after exposure to 0.5 M NaCl solution for 42 h at 40 °C.

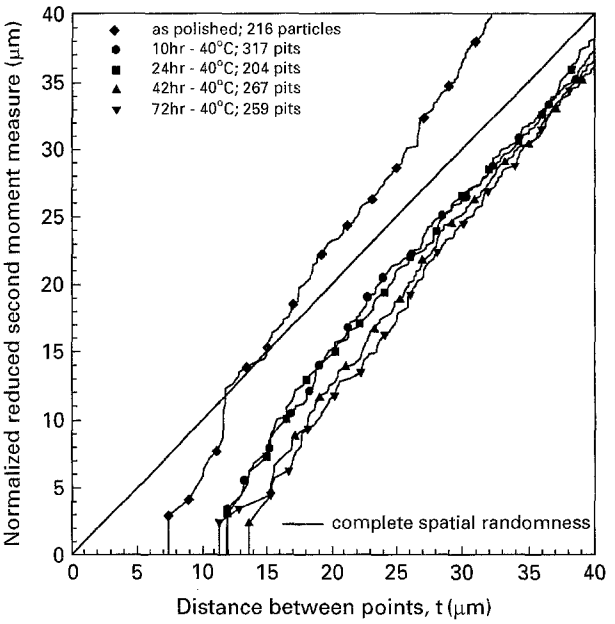

Figure 3 Estimation of $\hat{K}(t)$ for 2024-T3 aluminium alloy specimens before and after exposure to 0.5 M NaCl solution for 10, 24, 42, and 72 h at 40 °C.

that the estimated curves fall well below the diagonal for all values of t. This implies that the centroids of the pits tend to exhibit regularity as opposed to clustering. This observation is consistent with typical corrosion pitting because particles which are quite close coalesce fairly early in the corrosion process to form a larger pit. Thus, pitting tends to encompass clustered particles. On the other hand, the estimate of $\hat{K}(t)$ for the polished specimen is well above the diagonal for most of the range. Qualitatively, this implies that particles tend to be more clustered. The fact that the lower tail falls below the diagonal may be partially due to measurement limitations of the optical microscopy. Clearly, the character of the spatial statistics of particles prior to exposure to the environment is different from that for the corrosion pitting process. It is also manifest that neither the particle nor the pit data can be represented by the CSR model. Consequently, a more detailed model is warranted, and furthermore, it is expected that different models will be required for the particle and the pit data.

As mentioned above, the class of models which is assumed for this application consists of the hard-core models. This class is for spatial data whose features do not intersect. It is granted that this assumption is not in accord with the actual evolution process; however, at any fixed time, it is not likely that coalescence of pits could be observed. Thus, for typical data a hard-core model should be close to observations. The hard-core model considered herein is a modification of Diggle's simple sequential inhibition (SSI) model [6]. The modification consists of replacing the geometrical objects which have a fixed size and shape, i.e. in Diggle's model the objects are discs of constant radius, with objects which have a random size and shape. Henceforth, the proposed model will be referred to as the random sequential inhibition (RSI) model.

Because the primary method of analysis for the RSI model is by simulation, the various aspects of the procedure must be described. There are two key components needed for the construction of the RSI model. The first is the selection of the shape and size of the objects used to construct this hard-core model. The initial attempt is to assume that circular discs with random radii are reasonable approximations for both types of features. The cumulative distribution functions (cdf) for the radii are estimated by using the measured attributes of the features. The attributes considered included the shortest dimension, the longest dimension, an equivalent radius calculated from the area by assuming a circular shape, and an average of the shortest and the longest dimensions. For the data from the corrosion specimens, the best fit to the proposed model was obtained using the longest pit dimension. On the other hand, the best fit to the other model for the polished specimen data resulted in using the average of the longest and the shortest particle

dimensions. A two-parameter Frechet cdf was fit using the standard maximum likelihood estimation (MLE). The standard two-parameter Frechet cdf is

$$F(x) = \exp[-(x/\beta)^{-\alpha}] \qquad x \geqslant 0,\ \alpha > 0,\ \beta > 0 \quad (8)$$

The Frechet cdf was selected because it is the cdf associated with the maxima of random variables (rvs) which is concentrated on the non-negative reals. A cdf based on the maxima of rvs is preferred because it is the largest features from which the most severe damage ensues. The data and the Frechet fits are included in Fig. 4 for all five of the specimens considered. The Frechet MLE estimates fit the data quite well. The estimates for the parameters in Equation 8, the mean μ, the coefficient of variation, cv, and the sample size, n, for the five specimens are contained in Table I. It should be noted that the scatter is quite large for all the specimens. This large scatter is an indication that the observations for the specimens subjected to the environment contain pits which are at all stages in the evolution of nucleation and growth. Also, the constituent particles have a wide range of sizes. It should be noted that the mean feature size is increasing as the exposure time to the environment increases.

The second key component in the RSI model required for the simulation is the distribution for the centres of the circles. This distribution is required for the description of the point pattern. However, this is not known, and furthermore, it cannot be estimated directly. This is the heart of the implicit estimation that is needed. the key aspect of the simulation is the placement algorithm for the centres. The centre of the first circle is uniformly distributed over the entire area, and its radius is simulated from its cdf given above. The centre of the second circle could be uniformly distributed over the entire area minus the area occupied by the first circle; however, this would not produce the desired regularity for the cavities. In order to avoid too much clustering, the centres of the simulated cavities are to be separated by at least a given distance, D. This distance D is estimated by the minimum of the distances between all the centroids in the data set. Thus, the centre of the second circle is uniformly distributed over the original area minus a circular region of radius equal to the maximum of the radius of the first circle and D. The radius of the second circle, again, is simulated from its cdf. If the two simulated circles do not intersect, then the algorithm continues. If the first and second circles intersect, or if the distance between the two centres is less than D, then the second circle is discarded entirely, i.e. the centre and its radius, and another circle is generated. The algorithm continues by randomly positioning each successive circle over the remaining available area such that it is disjointed from all of the previously positioned circles and its centre is at least D units from all of the previously simulated circles. Clearly, the number of circles needed to be generated in order to produce the required pattern of circles is greater than the observed number of cavities. The total number of simulated circles needed for the polished specimen is only about 17% greater than the actual number, but for the 72 h, 40 °C specimen it is about 74% greater. The construction algorithm ends when the number of circles satisfying the given conditions equals the number of observed features.

The estimates for $K(t)$ and $G(y)$ play a vital role in assessing the suitability of the assumed RSI model. The final task is the comparison of the simulated results with the observed data. The test of the suitability of the implicit model will be a 99% confidence interval for the estimates of $K(t)$ and $G(y)$. The 99% confidence intervals are estimated from 99 simulations of the proposed implicit model. For each simulation i, the estimates $\hat{K}_i(t)$ and $\hat{G}_i(y)$ are computed. The upper and lower confidence bounds are estimated by the maximum and minimum of $\hat{K}_i(t)$ and $\hat{G}_i(y)$ for $1 \leqslant i \leqslant 99$, respectively. Figs 5 and 6 show $\hat{K}(t)$ and $\hat{G}(y)$ with their confidence bounds, respectively, for the 42 h, 40 °C corrosion pit data. First, consider Fig. 5, and recall that t is the distance between events. Notice that $K(t)$ for the data lies within the estimated confidence interval over the entire interval for t, even though it is quite close for t near 15.5 μm. Based on this test, the proposed RSI model is suitable for modelling the pattern of the locations of corrosion pits. Furthermore, the confidence bounds are reasonably close which implies that the model is a relatively accurate representation of the data. Both the data and the RSI model cannot be characterized by the CSR model for nearest neighbours, because they deviate appreciably from the diagonal. The assumption that

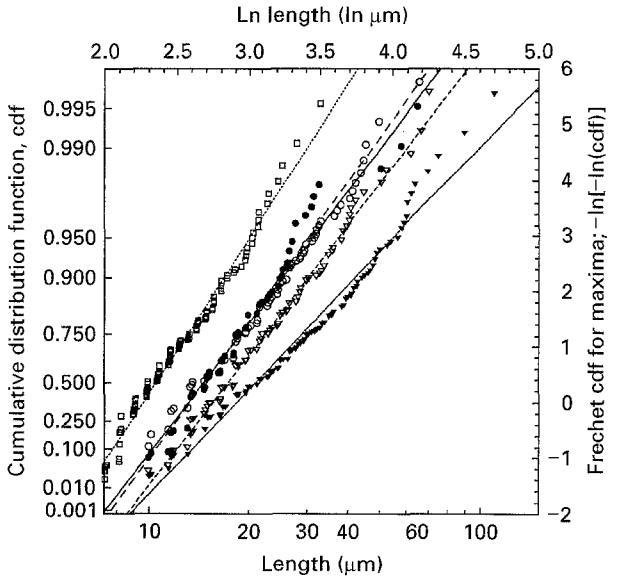

Figure 4 MLE estimation of the feature radii using a two-parameter Frechet cdf for 2024-T3 after exposure to 0.5 M NaCl solution. (□) Polished (typical); (○) 10 h, 40° C; (●) 24 h, 40° C; (▽) 42 h, 40° C; (▼) 72 h, 40 °C.

TABLE I Estimated parameter values for the observed features for the five specimens.

Specimen	n	$\hat{\alpha}$	$\hat{\beta}$(μm)	$\hat{\mu}$(μm)	$\hat{cv}$(%)
Polished	216	4.45	9.53	11.38	36.6
10 h, 40 °C	317	3.56	13.12	16.66	50.7
24 h, 40 °C	204	4.00	13.75	16.85	42.5
42 h, 40 °C	267	3.38	15.59	20.14	55.0
72 h, 40 °C	259	2.51	17.94	26.63	101.9

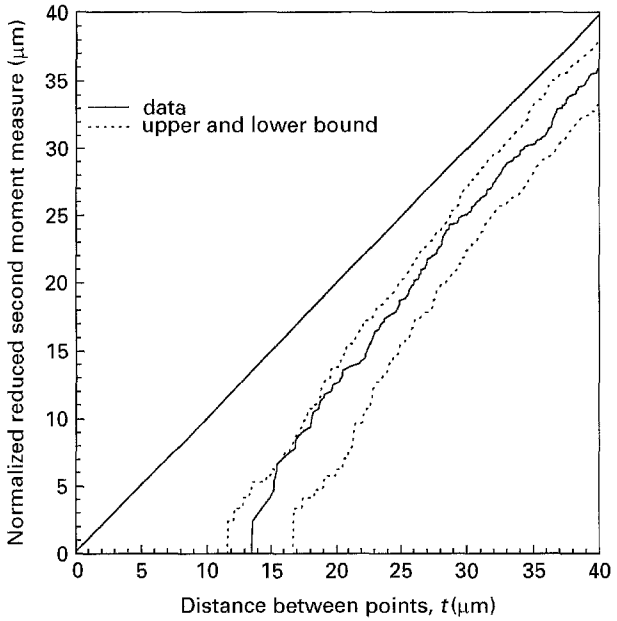

Figure 5 Estimation of $\hat{K}(t)$ with the 99% confidence bounds produced from the RSI model for the 2024-T3 aluminium alloy specimen subjected to 0.5 M NaCl solution for 42 h at 40 °C. Distance between features exceeds a minimum value; 267 features, 99 simulations.

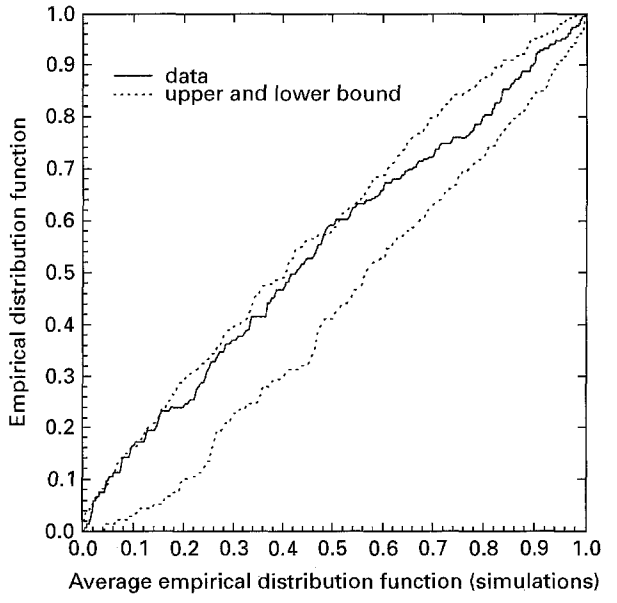

Figure 6 Estimation of $\hat{G}(y)$ with the 99% confidence bounds produced from the RSI model for the 2024-T3 aluminium alloy specimen subjected to 0.5 M NaCl solution for 42 h at 40 °C. 267 features, 99 simulations.

corrosion pits are uniformly distributed must be rejected.

An additional statistical test for the model is based on the edf $G(y)$ defined in Equation 5 and its estimator $\hat{G}(y)$ given in Equation 6. Fig. 6 is a plot for $\hat{G}(y)$ for the 42 h, 40 °C data and the corresponding upper and lower simulated 99% confidence bounds. These functions are plotted versus the average of the 99 edfs computed from the simulations. Close proximity of the curves to the main diagonal indicate that the data and the simulated model are reasonably matched. For the most part, $\hat{G}(y)$ for the data lies within the confidence bounds; however, for values in the interval (0.48, 0.52), $\hat{G}(y)$ exceeds the upper confidence bound by a very small fraction. This does not disqualify the proposed model for the cavity locations, but it does indicate that the model may be pushing the sensitivity and applicability for the data. Thus, this case clearly indicates the merit of more than one statistical test for assessing the validity of a model. Based on the statistics for both $K(t)$ and $G(y)$, the indication is that the proposed RSI model is acceptable for the corrosion pits.

Although the RSI models the locations of corrosion pits adequately, it is not appropriate for modelling the locations of constituent particles. As previously mentioned, the particles on the polished specimen tend to exhibit more clustering than the pits during the corrosion process. Therefore, a further refinement to the RSI model is required. The proposed model adds a generalization of a Poisson cluster process to the RSI model. The cluster process assumes an underlying parent process, with intensity ρ, from which offspring events are randomly generated. The parent process is assumed to be a homogeneous Poisson process; consequently, the parent events will be uniformly distributed over the area. Each parent produces a random number, S, of offspring events which are assumed to be independent and identically distributed according to a homogeneous Poisson process with intensity, μ. Again, at each parent, the offspring are uniformly distributed over the parental area. The intensity, λ, of the resulting cluster process is given by

$$\lambda = \rho\mu \tag{9}$$

This model therefore will be referred to as a cluster random sequential inhibition (CRSI) model.

The CRSI model will also rely on simulation for analysis, with the procedure as follows. Parent events are assumed to be circular discs with a deterministic, common radius. These discs are placed uniformly over the area, using the sequential algorithm previously described, such that no two discs intersect. The number and size of the parent discs are based on observations of the graph of the locations of the particle centroid data. These parameters are related, and they affect the resulting amount of clustering of the offspring events in the model. Once the parent locations have been determined, the simulated offspring events are positioned. The offspring events are assumed, as in the RSI model, to be circular discs of random size. The cdf for the radii is assumed to be a two-parameter Frechet distribution given in Equation 8, and it is estimated from the average of the longest and shortest dimensions of the observed particles. The parameter estimates are contained in Table I. The offspring locations are determined as follows. A parent is selected at random, then an offspring event, with random radius, is located uniformly over the parent area. To place the second offspring, a parent is again selected at random. The offspring then is placed uniformly on the remaining area, using the RSI algorithm previously described. That is, the simulation ensures that no two offspring intersect and that each offspring is separated from all other offspring by a minimum distance, D,

which is estimated by the minimum distance between all of the centroids in the data. This algorithm continues until the number of offspring placed is equal to the number of observed particles.

To illustrate the CRSI model, Fig. 7 is a schematic representation of one simulation for which there are 75 parents each with a radius of 35 μm. The maximum number of offspring per parent is 6. As expected, the parents, shown as the light grey circles, have from 0–6 offspring, which clearly causes some degree of clustering. The positioning of the offspring are based on the locations of their centroids so that a fraction of the offspring area may lie outside the area of the parent. It should be noted that the parent process and the offspring process are both hard-core processes as the above algorithm dictates.

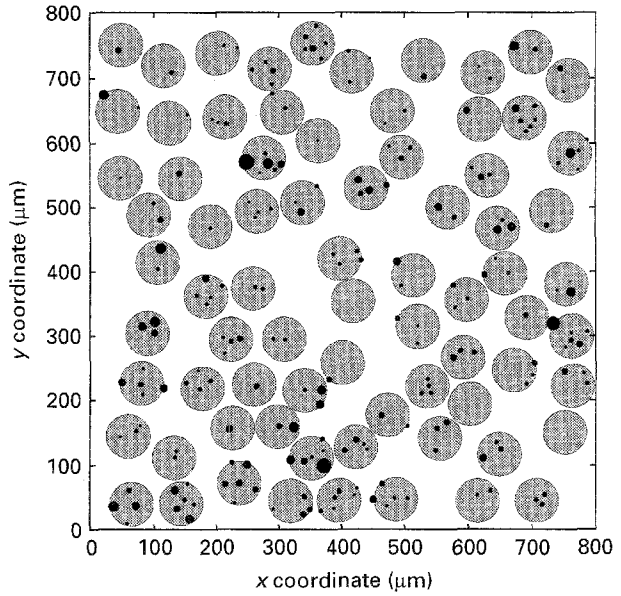

Figure 7 Schematic representation of a simulation of the cluster random sequential inhibition (CRSI) model.

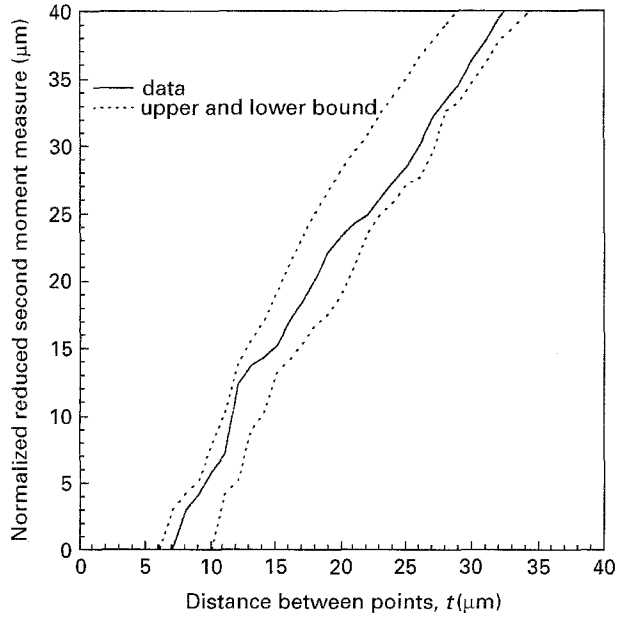

Figure 8 Estimation of $\hat{K}(t)$ with the 99% confidence bounds produced from the CRSI model for the 2024-T3 aluminium alloy specimen before exposure to 0.5 M NaCl solution. 99 simulations; as-polished; 216 features; 75 parents; parent radius = 35 μm; maximum number of offspring per parent = 6; Frechet cdf for offspring radius, ρ.

Comparison of the simulated results with the observed particle data will again involve the statistics, $\hat{K}(t)$ and $\hat{G}(y)$. Fig. 8 shows $\hat{K}(t)$ for the particle data, with the upper and lower 99% confidence bounds from the proposed model. Note that $\hat{K}(t)$ for the data lies within the estimated confidence interval over the entire range of t, which indicates that the CRSI model is acceptable for modelling the locations of the particles. The comparison based on the edf $G(y)$ lends additional support to the applicability of the proposed model. Consequently, the proposed CRSI model is a suitable choice for modelling the locations of constituent particles.

5. Conclusion

The particles and the corrosion pits after exposure to the environment in a 2024-T3 aluminium alloy are characterized by random locations, random shapes, and random sizes. The random spacing of the particles and the corrosion pits can have a significant impact on the integrity of 2024-T3, at least for certain environmental conditions. Reliability computations must include the rather complex random geometrical properties of particles and corrosion pits. Corrosion pits tend to be regularly spaced, whereas particles tend to be more clustered. Pragmatically, the RSI hard-core model composed of circles with a random diameter and a random location is very reasonable as an approximation for the spatial pattern of the centroids of corrosion pits, based on the conditions considered. The RSI model is not applicable for modelling the spatial statistics of the constituent particles; however, the CRSI, which incorporates an underlying random parent process, does provide reasonable approximations for the spatial pattern of the particle centroids. The simplicity of the models has two benefits. The models are computationally tractable, at least with standard simulation techniques, and they can be employed for a variety of conditions. The qualitative and quantitative accuracy of the models are manifest from observations of the spatial patterns and the relatively tight confidence bounds on the computed statistics for $\hat{K}(t)$ and $\hat{G}(y)$. Further research, including modelling and experimentation, is needed to include the combined effects of time and temperature, for a given environment. This will be considered in the future.

Acknowledgements

This work was supported in part by the US Air Force Office of Scientific Research under Grant F49620-93-1-0426 and by the Federal Aviation Administration under Grant 92-G-0006.

References

1. B. D. RIPLEY, "Spatial Statistics" (Wiley, New York, 1981).
2. P. J. DIGGLE, "Statistical Analysis of Spatial Point Patterns" (Academic Press, New York, 1983).

3. D. STOYAN, W. S. KENDALL and J. MECKE, "Stochastic Geometry and Its Applications" (Wiley, New York. 1987).
4. N. A. C. CRESSIE, "Statistics for Spatial Data" (Wiley, New York, 1991).
5. G. S. CHEN, M. GAO and R. P. WEI, *Corrosion* **52** (1996) 8.
6. P. J. DIGGLE, J. BESAG and J. T. GLEAVES, *Biometrics* **32** (1976) 659.
7. P. J. DIGGLE and R. J. GRATTON, *J. R. Statis. Soc. B* **46** (1984) 193.
8. B. D. RIPLEY, *J. Appl. Prob.* **13** (1976) 225.
9. *Idem, J. R. Statist. Soc. B* **39** (1977) 172.

Received 8 August
and accepted 21 December 1995