Maintenance of Water Distribution Systems

Tim Watson, Dr Colin Christian
Department of Civil and Environmental Engineering
Dr Andrew Mason, Dr Murray Smith
Department of Engineering Science
www.esc.auckland.ac.nz
The University of Auckland

Abstract

The efficient long term management of large-scale public funded assets is an area of growing importance. Ageing infrastructure, growth, and limited capital all result in the need for more robust and rigorous methodology to prioritise rehabilitation and renewal decisions and more importantly, to forecast future expenditure requirements. The overall objective of this research is to develop a Bayesian-based decision support system that will facilitate the identification of efficient asset management policies. This paper focuses on areas of research relating to the long term management of water distribution systems and in particular, will present: 1) A general review of existing practice and research; 2) development and initial results for an object oriented discrete event simulation, and 3) proposed future research and development.

1 Introduction

A water distribution network can be described as a network of pipes, valves, and pumps (links), that supply or distribute water to end-users (nodes). It is the pipes however that make up the bulk of the capital cost of the system. Figure 1 shows an example of a water distribution network.

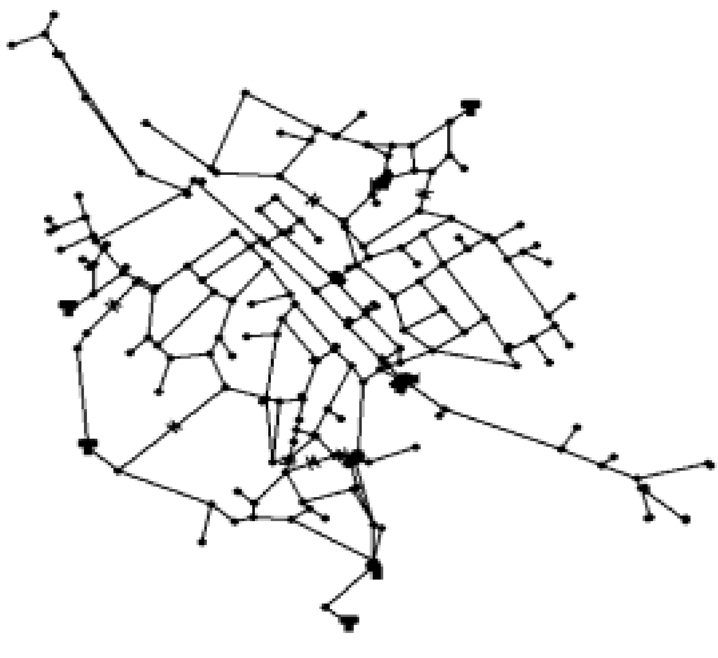

Figure 1. A Typical Water Distribution Network

It is well understood that water supply infrastructures throughout the world are deteriorating systems. This, coupled with further expansion and renewal, places ever-increasing financial stress on present levels of funding. Water rates in Auckland, for example, are predicted to increase 5-20% per year for the next 10-20 years just to keep pace with maintenance and renewal programmes. The water industry is therefore facing the problem of managing deteriorating networks in the most efficient manner possible to maintain existing and future levels of service.

Levels of service are normally based on predetermined criteria that are often listed by a water supplier as a ‘customer charter’ agreement. Additionally, there are statutory obligations that a supplier has to comply with. The main targets usually include, as a minimum, delivery of water at a minimum pressure and flow, limited downtime per year, fire fighting provisions, and water quality standards. In practice, due to the nature and design of the networks, primarily relating to built-in redundancy and storage, supplying minimum pressure is the overriding criterion that has to be achieved.

It is obvious that for a deteriorating water supply system to meet predetermined levels of service, maintenance actions need to occur. Maintenance is carried out to prevent system failures, as well as to restore the system function when a failure has occurred. The prime objective of maintenance is thus to maintain or improve the system reliability and operation regularity [1].

Maintenance typically consists of rehabilitation, repair and renewal. Most maintenance policies combine the making of renewal decisions after one or more failures with planned pipe replacements based on engineering judgement and knowledge of the system.

In order to enhance maintenance decisions it is therefore essential to improve the understanding of the deterioration process and the evolution of failures of water mains pipes and to develop appropriate predictive models that can assist in the decision making process. The resulting pipe break models can serve both as a diagnostic tool and an optimisation tool (e.g. for developing best replacement strategies), but also, when coupled with an economic assessment model, they become powerful tools for decision making for water managers [2].

2 Existing practice

There are several approaches and criteria used to determine replacement strategies for short and long term planning. The following section gives a typical example of the maintenance planning methodology of a large scale private water supplier [3].

Maintenance planning is split into the short-term (1-2 years) and long-term (5+ years). Long term planning tries to predict total projected annual costs based upon estimated asset lifetimes and an annual percentage renewal rate. This ensures that the average replacement age of a pipe is equal to the estimated lifetime. This process is reviewed every five years and included in the Strategic Asset Management Plan.

Short term planning is done on a yearly basis and included in the Annual Asset Management Plan. It focuses on what individual assets need to be replaced in keeping with the long term forecasted annual costs. A schedule of capital works, mainly pipe replacements, is catalogued and budgeted for the forthcoming year. Additionally, an emergency budget is allocated for pipes that need to be urgently replaced without waiting for the following year’s budget. This is estimated to be on average 5-10% of the scheduled annual budget.

The schedule of capital works is based primarily on estimations of the pipe break rates and age of pipes. This is done by using historical failure and asset information data, and more predominantly, engineering judgement and experience. Break rates for individual pipes, or perhaps groups of ‘similar’ pipes, are estimated and then used to predict a likely stream of costs over a future period, say for example ten years. If the expected repair costs over the ten years are greater than the expected replacement costs, then the pipe is replaced. An estimated growth rate in the pipe failures is sometimes applied where it is felt appropriate. Customer outage hours are also taken into consideration but can be difficult to include in the initial cost estimation process. This is therefore typically added as a secondary, independent criterion.

Repair is always actioned for a pipe failure, whether or not it is followed by replacement. This repair is necessary as any replacement work is likely to take a significant period of time, and so customers need to be informed before this occurs. Repair is performed by a pipe repair contractor who also records (very) limited details about the nature and extent of the failure. In rare cases, where the contractor thinks the state of the pipe is extremely poor, the pipe is scheduled for emergency repair. However, this is often not in the best interests of the contractor as it reduces future breakage rates.

A general result of the above maintenance approach is that the whole planning process becomes very subjective and ad-hoc. Rules of thumb tend to dominate replacement decisions rather than robust statistical analysis. Such rules relate maintenance decisions to pipe age, material, number of previous breaks, and location, amongst other factors. Additionally, long term yearly budgets may already have been decided before any yearly analysis has been undertaken.

3 Previous Research

The majority of existing research related to the modelling of pipe failures can be split into three main categories: survival analysis, aggregated regression models, and probabilistic predictive models.

Survival analysis focuses on the lifetime of a pipe and is predominantly used for long term financial planning. The pipe lifetime is treated as a random variable and a standard statistical distribution is then fitted to a collection of similar pipes. The pipe group can then be aged to assess what the likely future costs of replacement will be. Such analysis has been performed for several European water networks [4, 5] and North American networks [6].

Aggregated models group together pipes that have the same intrinsic properties and then use linear regression to establish a relationship between the age of the pipe and the number of failures. Shamir and Howard [7] proposed an exponential increase with time of the form

$$\lambda(t) = \lambda(t_0)e^{A(t-t_0)}$$

where $\lambda(t)$ is the number of failures/yr/1000 ft at time t, t_0 is the base year for analysis, and A is the growth rate coefficient. A power law function has also been used as the growth function [8].

A number of researchers have used multiple regression to extend the above exponential equation to relate the environmental and intrinsic properties of the pipe with the failure rate as a function of time [9, 10]. The environmental and intrinsic properties, or explanatory variables, that are most commonly used are the pipe material, diameter, soil type, pressure, and the nature of the surrounding land development.

Probabilistic predictive models try to predict the probability that a pipe will burst at a particular time. This can then be used to calculate the economic life of a pipe and therefore when it should be replaced. Andreo et al. [11] proposed the use of a Cox proportional hazard model [12] to relate the hazard function to a set of explanatory variables. The basic form of this model is

$$h(t:z) = h_0(t)e^{zb}$$

where, $h(t{:}z)$ is the failure rate (termed hazard function; see later), $h_0(t)$ is some unspecified baseline hazard function, z is a vector of explanatory variables (diameter, soil, etc.), and b is a vector of regression coefficients. Further research using a proportional hazard approach includes Gat and Eisenbeis [13] and Lei and Saegrov [14]. Both these authors used a Weibull hazard model to model the useful life of a pipe.

A very brief description of the main shortcomings in the above models is as follows. The survival analysis approach groups similar pipes and relies heavily on estimating the lifetime of the groups, which may itself be highly variable and dependent on the individual pipe characteristics. Aggregate regression models do not in general provide any information about individual pipes and the variability that may exist between individual pipes. Therefore, fitting standard distributions across a group may not be appropriate. Additionally, these models have relied on complete data sets for derivation. The proportional hazards models are more sophisticated but typically lack statistical rigour in their formulation and suffer from a reliance on complete long-term failure records.

In general, all models have a tendency to not fully recognise and model appropriately the fundamental statistical processes that are occurring. They try to fit models around the data, rather than the data into an underlying model. The next section presents an underlying statistical model for pipe failures.

4 Repairable versus Non-Repairable

A repairable system is a system that, when a failure occurs, can be restored to an operating condition by some repair process other than replacement of the entire system [15]. A water supply network can therefore be defined as a repairable system. If a pipe fails, it can either be replaced or repaired to restore the system to an operating state.

A nonrepairable system is one which is discarded and/or replaced after its first failure. A light bulb for example, is a nonrepairable system. Further to this obvious example, the definition of nonrepairable can be expanded to include economic considerations. For example, a cheap watch can be repaired if broken, but due to the cost of repair being greater than the cost of purchasing another one, it would probably be discarded and replaced. It can therefore be considered nonrepairable.

From a purely physical perspective, it is intuitively obvious that an individual water pipe can be considered to be a repairable item. Work can be done on the pipe to restore it to a functioning state, even if it only lasts for a short period of time. However, when considering the cost of failure, the situation is not quite so clear. If, for example, there is little redundancy in the network, a pipe failure might have significant consequences (repair and failure costs) greater than the replacement cost. It is therefore efficient to replace the pipe at some time period before or equal to failure depending on the failure/replacement cost ratio.

For a nonrepairable pipe, the lifetime, T, is a random variable. Repair is not considered and the component is replaced or rehabilitated in a way that restores the pipe to an "as good as new condition". It is assumed that the failure of one pipe does not affect the performance of a similar pipe located elsewhere. It is also assumed that if all the pipes are similar then the pipe lifetimes have the same distribution. These two assumptions can be combined into one statement that says the lifetimes are independent and identically distributed. This can be described statistically as a renewal process.

A renewal process can be modelled by fitting a standard statistical lifetime distribution. This can then be used to model the hazard rate. The hazard rate, h(t), is the conditional probability that a unit will fail in the time interval (t, t+Δt], given that the unit is functioning at time t. It can be written as:

$$h(t) = \frac{f(t)}{1-F(t)} = \lim_{\Delta t \to 0} \frac{P\left(t < T \le t + \Delta t \mid T > t\right)}{\Delta t}$$

where T is the random lifetime of the pipe. This can then be combined with the consequences of failure to calculate the risk of an asset for any given time as

$$Risk(t) = h(t) \times Consequences$$

If the consequences can be sensibly converted into monetary costs, then the risk becomes an expected cost of failure and can be compared to the expected cost of replacement. Financial risk-based techniques such as net present value can then be used to optimise the maintenance strategy and schedule appropriate replacement.

Fortunately, for a well-designed water supply system, the bulk of the pipes can be considered repairable as the cost of failure is small in comparison to replacement. There are three main reasons for this. Firstly, the environmental effects of a failure are minimal; secondly, there are storage reservoirs within the network; and thirdly, networks typically contain redundancy and diversity.

When a pipe failure occurs, a small part of the pipe can be repaired or possibly replaced, to restore the pipe to its functioning state, without replacing the entire pipe. The pipe may be repaired several times before being replaced. Therefore, pipes that are considered repairable cannot be modelled by the conventional fitting of a statistical lifetime distribution, as successive failures are firstly not identically distributed and secondly not independent. The succession of pipe failures can, however, be modelled using a counting process.

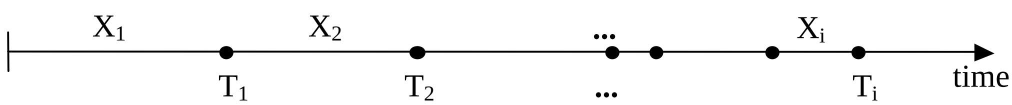

Figure 2. Time Dot Plot of a Counting Process

N(t) can be defined to be a random variable denoting the number of failures for a given pipe in the time interval (0,t]. The process $\{N(t), t \geq 0\}$ is called a stochastic process, or more specifically a counting process. A counting process is a stochastic model that describes the occurrence of events in time. These occurrences, in this case failures, can be thought of as points on the time axis and are commonly displayed graphically as in Figure 2. The recorded failure times are denoted $T_1 < T_2 < \ldots < T_i$ and the corresponding time between failures $X_1, X_2, \ldots, X_i$. If the times between failures tend to get shorter with age the item is said to be deteriorating. Alternatively, if the times between failures are increasing then the item is improving.

The expected number of failures from time 0 through to time t is known as the mean function, *Λ(t)*, and can be written as:

$$\Lambda(t) = E(N(t))$$

The intensity function, λ(t), of a counting process at time t is defined to be:

$$\lambda(t) = \frac{d}{dt}\Lambda(t) = \frac{d}{dt}E(N(t)) = \lim_{\Delta t \to 0}\frac{E(N(t+\Delta t) - N(t))}{\Delta t}$$

The intensity may be regarded as the mean number of failures, N(t), per unit time. In other words, this describes the rate at which breaks are occurring. The intensity function is therefore the most important measure of a pipe's reliability. Many repairable systems, including pipes, typically have a "bathtub" shaped intensity function. This is shown in Figure 3.

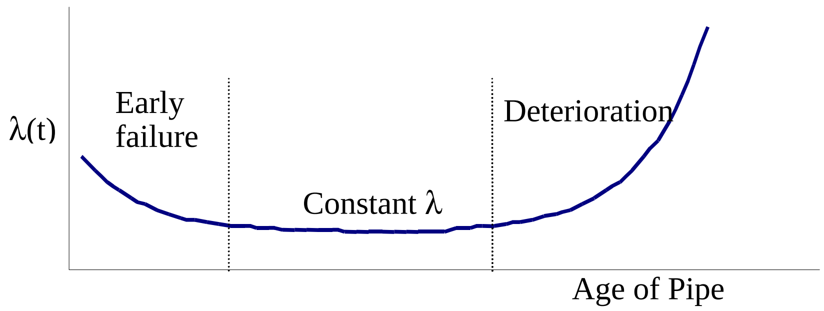

Figure 3. Bathtub Intensity Function

When a pipe is newly installed the intensity can be high and failures frequent. This can be described as a settling in period, possibly due to construction practices. After early faults have settled down, the intensity will be smaller and remain relatively constant for long periods of its useful life. Then as the pipe ages, the intensity will begin to increase and the pipe start deteriorating. This is the period of most interest, as eventually the intensity will exceed a certain level, and it will become cost efficient to replace the pipe.

This deterioration phase can be modelled by a Nonhomogeneous Poisson Process (NHPP).The NHPP allows the intensity to vary with time. This implies that the times between failures are neither independent nor identically distributed. The NHPP also assumes minimal repair with negligible repair times. Minimal repair means that a failed pipe is restored back to its just-functioning state. The likelihood of the pipe failing is the same immediately before and after a failure. A minimal repair restores the pipe to an "as bad as old" condition. This assumes that a repair will have no effect on the intensity or the pipe failure rate. Negligible repair time is considered acceptable when comparing a pipe's lifetime measured in years with repair times counted in hours.

There are several models that can be used to describe the intensity of a NHPP. Common ones include the power law model ($\lambda(t) = at^{b-1}$), the exponential model ($\lambda(t) = ae^{bt}$), and the proportional hazards model ($\lambda(t) = \lambda_0(t)e^{bz}$) [12]. Once the intensity function has been estimated the number of failures in a given time period is Poisson distributed:

$$P(N(t) = n) = \frac{\Lambda(t)^n}{n!}e^{-\Lambda(t)}$$

5 Proposed Approach

Many factors influence maintenance decisions. These may be economic, environmental, operational, and even political. As well, it is very difficult to get water network managers to adopt a methodology in which they have little understanding and faith. This is particularly important as engineering judgement plays such an important role in the process. The objective of this research is not to provide water managers with a seemingly complicated statistical model that provides a single value of the useful life of a pipe, but instead offer a range of scenarios and policy outcomes that are useful for making maintenance decisions. Additionally, we want to capture all existing knowledge, both anecdotal and historical, for the water network being modelled.

We are developing a decision support system to assist managers in their network maintenance operations. There are two main components to our proposed system: a Bayesian statistical model and a discrete event simulation. These are discussed in more detail in the following sections.

6 Bayesian Statistical Modelling

The use of databases, and in particular geographic information systems, for managing infrastructure has only become common in New Zealand within the last ten years or so. The often poor quality of earlier paper records means it is very rare for water utilities to have an entire history of failures. Recorded failure data is more often than not limited to recent periods of time. A lot of the data that has been recorded is of dubious quality, being recorded or inputted wrongly, and quite often not in the right form or format. Research in European countries has identified that a smaller amount of more accurate data can lead to better results than more complete, but uncertain data [13]. The unreliable nature of the data creates several problems in modelling the deterioration process of pipes. The most obvious one is that with a lack of failure history it becomes very difficult to estimate the failure rates of pipes. This has been a major shortcoming in previous research and has resulted in the dominance of engineering judgement in the decision making process, as discussed previously.

Even when there is reliable failure history, the variability of pipe failures is large. This means that even though two pipes may have exactly the same failure rate, the number of observed failures can be quite different. A natural estimator of the failure rate of a pipe can be calculated by dividing the number of failures by the length of the pipe and the observed time period. Basing maintenance on the observed historical number of failures can perhaps incorrectly bias the decisions in favour of higher observed breaking pipes.

The problem therefore, is how to estimate the failure rate given limited reliable data and relying predominantly on engineering knowledge. One methodology that is well suited to this type of problem is Bayesian statistical modelling. A Bayesian model can combine engineering knowledge, in the form of our beliefs about the failure rate (prior distribution), with the data at hand to provide a formal estimate of the likely breakage rate distribution (a posterior distribution). This can be written as

$$P(\lambda \mid data) \quad \propto \quad P(\lambda) P(data \mid \lambda)$$

To further utilise existing data, we have constructed a hierarchical Bayesian model. Hierarchical models aim to combine the information from various sources of data while exploiting assumed similarity between parameters. In this model we assume that the underlying failure rates λ_i of each pipe i are drawn from the same prior distribution. This pooling of the data greatly improves the precision of the estimates of λ_i and can be seen as a compromise between the assumptions that all pipes are identical and therefore have the same failure rate, and that all pipes are different, in which case all pipes are treated separately using only data from pipe i to estimate λ_i. In this model, a pipe failure on one side of a network may provide knowledge about the failure rate of a similar pipe on the opposite side of the network. Similarly, information about a pipe failing due to a highly corrosive soil can be used to update our estimates of λ_i for all pipes in a similar soil type.

A constant break-rate model utilising this approach has been has been formulated based on the proportional hazards model. This can be used to provide insights into the factors contributing to failure, such as diameter, installation date, soil type, etc. Although theoretically tractable, the constant failure rate in this model means it has limited use for medium and long term maintenance planning. We are working to remove this limitation.

7 Simulation

An object oriented discrete event simulation has been developed to test various network maintenance policies. The simulation consists primarily of a Network Manager, a Pipe Manager, a Statistics Manager, and a collection of pipes and nodes that make up the network. Events are randomly generated based on some predetermined failure rate and placed on an event list. Events can be either a pipe failure, pipe repair, or pipe replacement, and each event has an associated event time and cost. When each event becomes active it is passed to the Network Manager who decides what action should be taken. If for example a pipe is to be replaced, then the Network Manager instructs the Pipe Manager to replace the pipe, which in turn schedules a pipe replacement event and updates the pipe details.

The main advantage of this simulation approach is that it is very transparent and understandable for real Network Managers, increasing their faith in and ultimately use of our new approaches. As well, it can be easily adapted to handle a wide range of criteria and can output any results that may be of interest. Providing valuable insights into the decision making process and their effects is particularly important when the costs are not all monetary and a lot of the decisions are based on rules of thumb. The simulation can be used to find the optimal parameters given a predefined policy, or mixture of policies. Simple examples of policies that are commonly used include:

- Replace after ‘k’ breaks
- Replace when the break rate exceeds ‘b’
- Replace pipe after ‘y’ years
- Spend ‘m’ amount of money per year

To test the simulation, a ‘replace after 3 breaks’ policy was run for a simulated 1,000 years over a network consisting of identical pipes. A power function model using the same fixed shape and scale parameters for each pipe was used to model the pipe breakages. The results of this simulation run can be seen in Figure 4. A theoretical density can be calculated for the age of the pipe, R, at the k^{th} break

$$f_R(t) = -\frac{d}{dt}P(R > t) = \lambda(t)\frac{(\Lambda(t))^{k-1}}{\Gamma(k)}e^{-\Lambda(t)},$$

in terms of the break rate $\lambda(t)$ and the mean function $\Lambda(t)$. Using the power law model these are given by

$$\lambda(t) = abt^{b-1},$$
$$\Lambda(t) = \int_0^t \lambda(u)du = at^b.$$

These theoretical results are shown in the figure, and confirm that, as expected, the simulated results closely match those predicted by the theoretical model. Further experiments are planned for the simulation to explore more complicated underlying models and management policies.

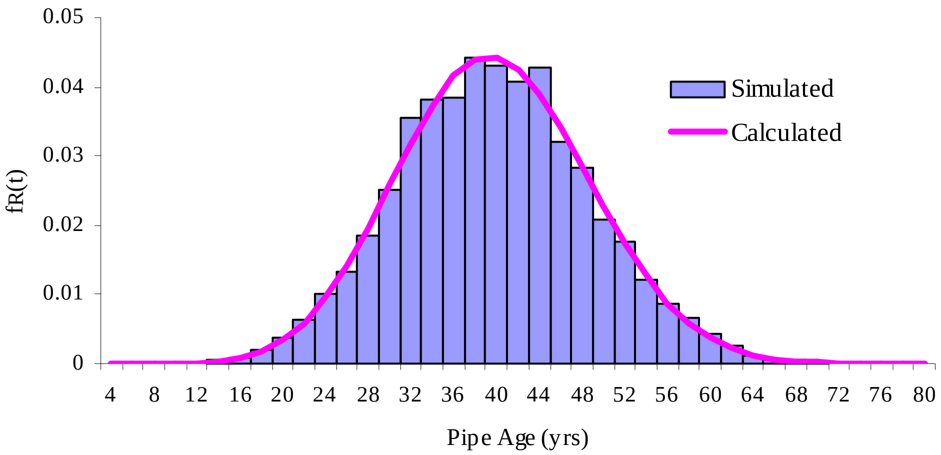

Figure 4. Three break policy using a power function model for pipe breakage rates.

8 Continuing and Further work

Further work needs to be done on the Bayesian model to develop a time varying growth model that can be used to predict future pipe failure evolution. This can then be used to predict medium and long-term maintenance and budgetary requirements. Additionally, work needs to be done to formally capture engineering knowledge as prior information.

Once a complete Bayesian model has been formulated it can then be combined with the simulation to better predict future expected costs. Furthermore, the simulation can be used to assess the value of information and the time period in which results can be used with a high degree of confidence.

References

[1] Hoyland, A. and M. Rausand, *System Reliability Theory*. Wiley series in probability and statistics. 1994, New York: Wiley & Sons, Inc.

[2] O'Day, D., *Organizing and Analyzing Leak and Break Data for Making Main Replacement Decisions.* Journal of the American Water Works Association, 1982. 74(November): p. 588-594.

[3] MetroWater, *Informal Maintenance Policy Discussions*, 2001.

[4] Herz, R.K., *Ageing processes and rehabilitation needs of drinking water.* Journal Water SRT, 1996. 45(5): p. 221-231.

[5] Herz, R.K., *Exploring rehabilitation needs and strategies for water distribution networks.* Journal Water SRT, 1998. 47(6): p. 275-283.

[6] Kleiner, Y. and B. Rajani, *Using limited data to assess future needs.* Journal of the American Water Works Association, 1999. 91(7): p. 47-61.

[7] Shamir, U. and C.D.D. Howard, *An analytical approach to scheduling pipe replacement.* Journal American Water Works Association, 1979. 71: p. 248-258.

[8] Tsui, E. and G. Judd, *Statistical Modelling of Water Main Failures*. 1991, Urban Water Research Association of Australia: Sydney.

[9] Clark, R., C. Stafford, and J. Goodrich, *Water Distribution Systems: A Spatial and Cost Evaluation.* Journal of the Water Resources Planning & Management Division, 1982. 108(October): p. 243-256.

[10] Walski, T. and A. Pelliccia, *Economic Analysis of Water Main Breaks.* Journal of the American Water Works Association, 1982. 74(March): p. 140-147.

[11] Andreou, S., D. Marks, and R. Clark, *A New Methodology for Modelling Break Failure Patterns in Deteriorating Water Distribution Systems: Theory.* Journal of Advanced Water Resources, 1987. 10(March): p. 2-10.

[12] Cox, D.R., *Regression Models and Life Tables.* Journal of the Royal Statistical Society, 1972. 34(Series B): p. 187-220.

[13] Gat, Y. and P. Eisenbeis, *Using maintenance records to forecast failures in water networks.* Urban Water, 2000: p. 173-181.

[14] Lei, J. and S. Saegrov, *Statistical Approach for Describing Failures and Lifetime of Water Mains.* Water Science and Technology, 1998. 38(6): p. 209-217.

[15] Basu, A.P. and S.E. Rigdon, *Statistical Methods for the Reliability of Repairable Systems*. Wiley series in probability and statistics. 2000, New York: Wiley & Sons, Inc.