A General-Purpose Operations Cost Model to Support Proactive Network Maintenance and More

Operations Value Model

A Technical Paper prepared for SCTE•ISBE by

Jason Rupe, Ph.D.
Principal Architect
CableLabs®
858 Coal Creek Circle
303.661.3332
j.rupe@cablelabs.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Introduction

This paper presents a flexible but basic set of Markov and math models that allow estimation of availability, reliability, and operations costs for repair operations under various levels of plant degradation. As a result, these models can address several use cases that can help operators make decisions about how and whether to address certain operations problems.

We are motivated by several observations.

- Proactive network maintenance (PNM) is an advantage given to cable operators for managing operations expenses to target service reliability which results from the underlying technology, usually DOCSIS® networking. But many operators do not take full advantage of it.
- Addressing the root cause of a problem takes much more technician time, so comes at a cost; but there is a benefit to the overall plant condition as well. It is important to examine this tradeoff, like so many other tradeoffs related to operations costs.
- Operators can be under pressure to improve services, which drives deployment of new technologies; but it is difficult to know where and when to make these transitions, and operations impacts are difficult to quantify.

In this paper, we explain the models in depth, define a few use cases for them, run these models using reference parameters to explore some possible generalizations useful to all operators, explain some questions that operators can answer with these models, and relate some maintenance optimization knowledge to the results.

Hopefully, this paper will entice operators into examining more operations improvement possibilities and enlist CableLabs to help explore some of them.

Model Description

There are two major models contributed in this paper: an 11-state repair model, and a degradation birth-death model. Both are Markov models, which requires transition time distributions to be exponential (though this can be relaxed through extensions). We add a small amount of mathematical modeling and software implementation of these models to combine them to address many use cases in operations. For readers not familiar with Markov modeling methods, there are many excellent references including Taylor and Karlan’s “An Introduction to Stochastic Modeling” [1].

The 11-state repair model describes the repair state for a section of cable plant that is in a static condition. With these 11 states, we can model both proactive and reactive repair, just proactive or reactive repair, or two parallel repair processes of any type at once. By exploiting features of this model, we can estimate the cost per unit time of an operations process it models. Therefore, we can use this model to compare operations solutions based on their effect on repair processing times, repair effectiveness, impact on customers, changes to costs, work handling, and other operations features. The 11-state model is shown in Figure 1.

Decoupled from the repair model, we have a degradation model based on a simple birth-death process. While the parameters that control the degradation and improvement of the condition of the plant can be modeled independently, we apply parameters to the repair model to better link them so that repairs that find and remove impairments will improve the plant condition, and causes of degradation that arrive externally will result in more repair work. A simple birth-death model is depicted in Figure 2.

When combined, the models can be viewed as multiple 11-state repair models, one for each number of impairments in the cable plant being modeled. Degradations shift the model to a set of repair states associated with a higher rate of failure; likewise, a repair that removes an impairment (failure cause) will improve the condition of the cable plant, and be reflected in the model with a move to a set of repair states associated with a lower failure rate. A view of how the 11-state repair model combines with the birth-death degradation model is shown in Figure 3. Note that above state w_0, where there is no degradation of the plant yet, there are no proactive repair steps shown (in green) because there is no proactive work to do. Also, Figure 4 shows a close up of the first two degradation levels with repair states and transition rates. The proactive states above state w_0 are shown translucent as they will not exist in most implementations.

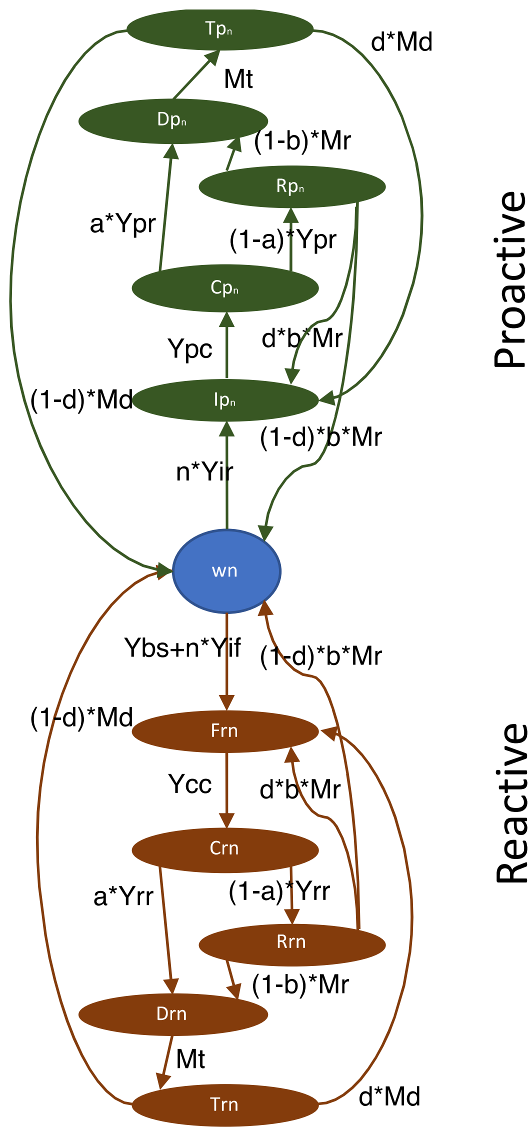

Figure 1 – A depiction of the 11-state repair model with proactive (green) and reactive (brown) states and transition rates shown.

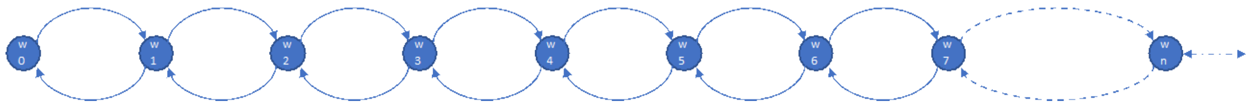

Figure 2 – A depiction of a birth-death process diagram as applied to the cable plant degradation model.

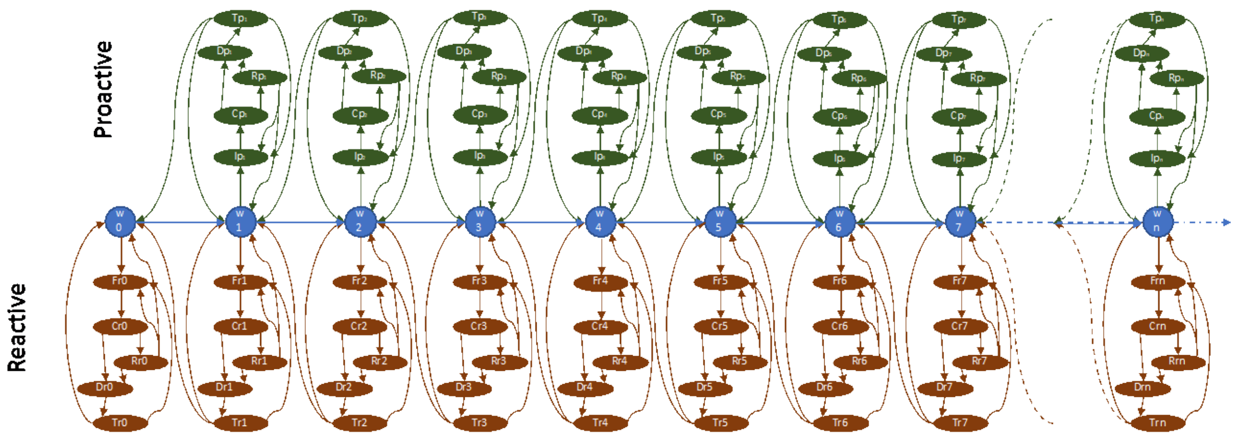

Figure 3 – A depiction of the 11-state repair model combined with the birth-death cable plant degradation model, showing the repair states at each degradation level, and how repair impacts the degradation.

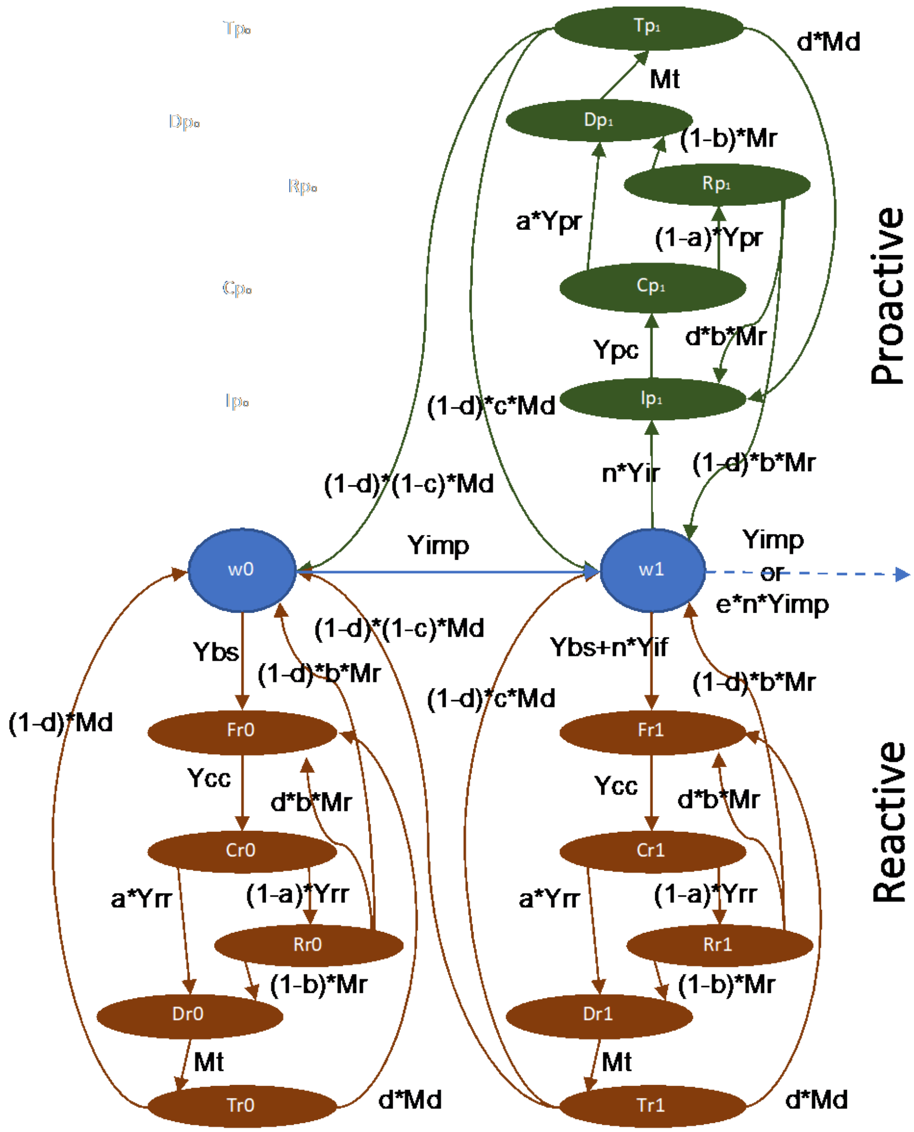

Figure 4 – A depiction of the first two degradation levels with the repair states and transition rates shown.

1. More About the 11-State Repair Model

The 11-state model is generally described as follows. It models the repair state of a section of plant serving a large enough number of customers to capture all customers potentially impacted by the largest such failure group, and small enough that it is unlikely that two failures would be addressed at the same time. About 1000 customers might be an appropriate group, for example. A key assumption to point out here is that the group of customers can experience one repair at a time, either due to proactive response or reactive response (depending on the model conditions), so we assume that failures and plant problems are rare enough that they get repaired before a second one happens. This seems reasonable for most cases, as it is for most maintenance models.

The model can be used to describe all failures for the section of plant, or a subset of the failure types. For example, only a set of failures that can be captured with a proactive system might be modeled for easy comparison. Likewise, a model of all possible failure modes would be created so that a degradation over time could be studied.

Because the repair model is small, and we expect to examine it over longer periods of time, we are mostly interested in the steady-state behavior of the repair system; we will often solve for the steady-state distribution of the 11-state repair model.

2. More About the Degradation Model

For the degradation model, if we assume there is an upper limit to the degradation that a reasonable network can sustain, finite state birth-death processes will work for the larger degradation model.

We can connect the repair model by determining the state probabilities and rates of not correcting the causes of the failures. Impairments represent degradation and can arrive according to an independent Markov process. Degradation should be slow, and not likely to reach a steady-state condition, so we need to study the transitive states of the model, not the steady-state. We will need to use a method called uniformization to solve this degradation model for various times over the lifetime of the cable plant. Because the model is a birth-death model, it solves sufficiently fast for reasonable time horizons.

A key assumption to point out here is that the group of customers can experience one repair at a time, either due to proactive response or reactive response (depending on the model conditions), so we assume that failures and plant problems are rare enough that they get repaired before a second one happens. This seems reasonable for most cases, as it is for most maintenance models.

3. Parameters

In the models described here, the following notation is necessary.

3.1. Definitions

- Impairment: This is a problem in the plant that will, after time and degradation, result in a failure that impacts one or more customers. These include loose connectors, nicks in the cable sheath, or a cable that swings in the wind, all of which eventually will result in a problem that worsens and needs to be corrected. Ideally, when corrected, the plant will be in better condition; but often times the plant gets minimally repaired, and so the impairment can result in another failure later.
- Shock: This is a type of event that results in immediate failure, with no proactive opportunity. Causes include storms, a shovel cutting through a drop, lightning, sabotage, or accidental damage. In these cases, the plant is damaged, and the damage immediately results in one or more customers having their service adversely impacted, so a repair is necessary.
- Remote repair: Without dispatching a technician, the problem impacting service is resolved. Perhaps this is done by a setting change in the system, or instructions with a customer. A remote repair attempt could be made but result in the need for a dispatch; this would be an unsuccessful remote repair attempt but is the same as a triage.
- Dispatch repair: This is a truck roll to send a technician out to locate and troubleshoot a problem to repair it. The repair may result in a minimal repair, where service is restored, but the plant is not improved. Sometimes this minimal repair actually fixes nothing, but removes the symptom that the customer experiences, such as placing a filter on the line to block certain frequencies. Or the repair may result in the cause of the problem being removed and the plant condition improved as a result. Improving the condition of the plant reduces the number of impairments in the plant, and thus the net rate of failures over time.

In the subsections that follow, we explain the model details further in terms of transition rates, states, and additional parameters. Figures 5 and 6 show verbose states and transitions for reactive and proactive repair process, respectively.

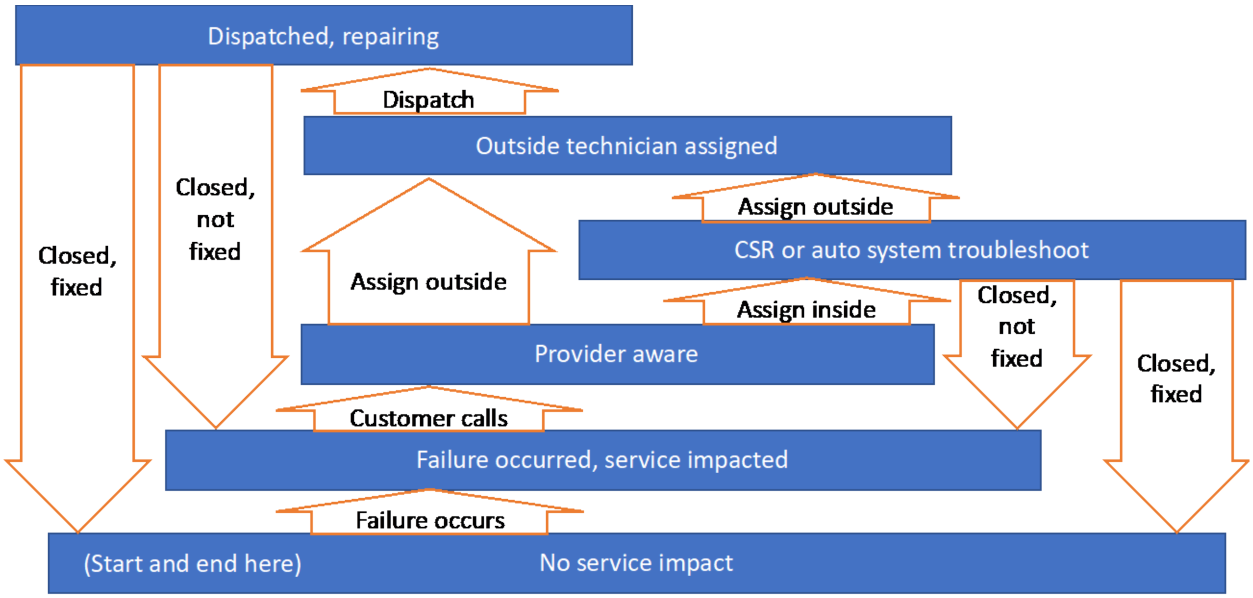

Figure 5 – A verbose depiction of the reactive lobe of the repair model, with the common state at the bottom, and reactive repair states above.

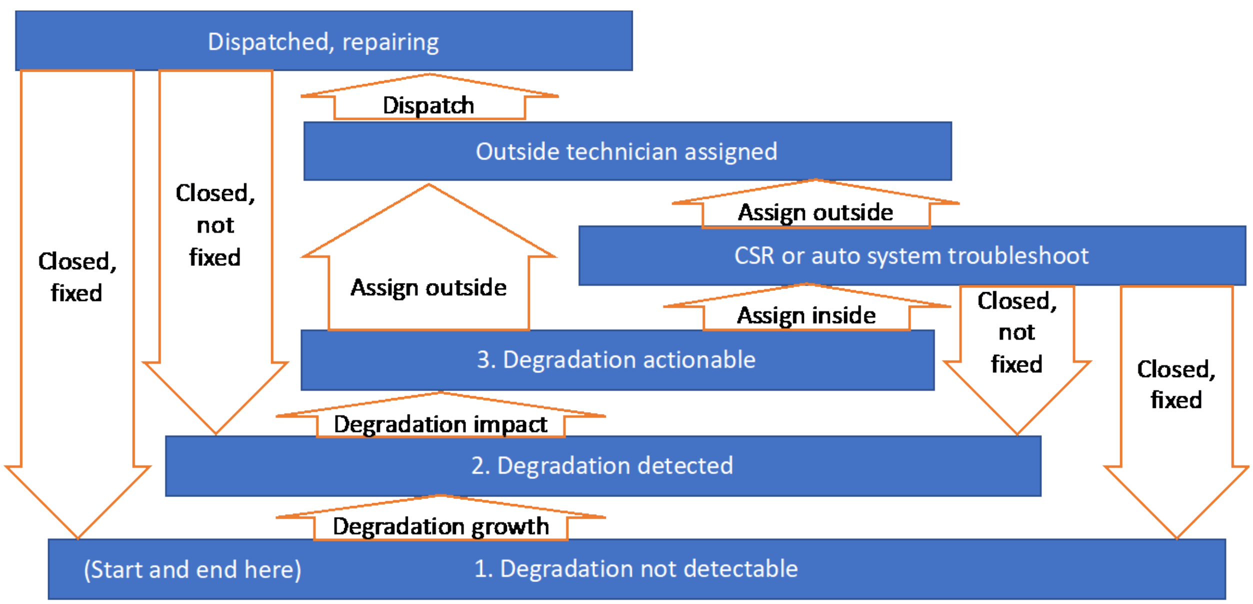

Figure 6 – A verbose depiction of the proactive lobe of the repair model, with the common state at the bottom, and proactive repair states above.

3.2. Transition Rates

Each of these parameters is described as a rate (events per unit time) because the model requires rates for input. But we can estimate the expected time each event type takes and estimate the rate as the reciprocal of the expected time.

- Y_{imp} = impairment arrival rate. This rate is not a part of the 11-state repair model but is used in the degradation model. The cable plant is exposed to the elements, and subject to damage. When damaged or degraded, there may not be immediate impact to service. But after time, the degradation becomes severe enough to lead to service impact. This measurement is the number of plant impairments (not service impairments) that arrive to a unit of plant in a unit of time. For example, for a section of plant, say a neighborhood, or all customers served downstream from a given amplifier, impairments might arrive due to damage or storms or erosion in the elements at a rate of one per month. While we know this is nearly impossible to tell, we can get at it another way as described next.

  Instead of the impairment arrival rate and impairment failure rate, we ask for the failure rate of a unit of plant, and the breakdown of the types of causes to the best of our knowledge. If we can determine the failure rate for a unit of plant (number of failures per unit of time) over a reasonably short period of time (say a year or less), and we can further determine the number of these failures that occur due to shocks versus degradations (by inspection and reporting from the technicians or center persons who work and close the ticket, or from the disposition codes or closure codes or comments), then we can approximate the factors needed for the model.
- Y_{bs} = base failure rate for shocks. The cable plant is subject to shocks that lead to failures. A cable modem (CM) can fail, an amplifier can fail, a cable can be cut, and so forth. Such plant failures are immediate, and different from degradations which can be detected proactively. The base failure rate from shocks can't be reduced without impacting external causes or changing the underlying technology such that you change the failure modes and failure rates of the external (extrinsic) shocks.
- Y_{if} = impairment failure rate. Plant impairments eventually result in service failures that impact customers. It is difficult to know in many cases how long an impairment has existed in the cable plant before it results in a failure. Therefore, we ask for other information that allows us to model this rate instead.
- Y_{cc} = customer call rate. This is the rate at which a customer, whose service is impaired, will call in the problem. The rate is events per unit time, so usually this rate is the reciprocal of the average time between when service is impaired to when they notice that, and then call in for service support.
- Y_{rr} = response to reactive call rate. This is the rate at which, once a customer calls in for help, the operator responds for action, either through an immediate dispatch, or by addressing the problem in a call or service center. The average time from when a customer calls to when the operator acts to resolve the problem is the reciprocal, and sufficient for this parameter.
- M_t = dispatch rate. Once a problem is decided to need action, if not resolved in a call or service center, and it is decided a dispatch or truck roll is required, this is the rate at which the dispatch is done. Sufficiently, we can use the time between when a dispatch is decided to when the dispatch is initiated.
- M_d = dispatch repair rate. Once dispatched, the average time to drive to the area, locate the trouble, and return the services to fully functional is sufficient for this rate parameter.
- M_r = remote repair rate. Once a customer calls in a problem and is connected to a call or service center, if the center resolves the problem, this is the rate at which they resolve the problem. Again, the average time to resolve a problem from the service center, without a dispatch, is

sufficient for this parameter. This estimate can include the center handling time for dispatched repairs too. Queue time should be included in any case.

- Y_{ir} = impairment reveal rate. Once an impairment is present in the network, it takes time before it impacts the radio frequency (RF) signal or other behavior of service, but well before the customer will notice it. This is where the sensitivity of your proactive monitoring and operations shows its impact. The average time between when an impairment exists in the network to when it can be detected using whatever methods you have in place for proactive detection will suffice for this parameter.
- Y_{pc} = proactive alarm rate. Once an impairment is detected by a proactive method, it will exist in the network and impact service until it rises to the level requiring action. This is where the operations method, culture, and practices are revealed best. Provide the average time from when a proactive issue is detected to when it gets bad enough to be addressed by your program, and the rate will be the reciprocal in most cases.
- Y_{pr} = response to proactive rate. Once action is warranted on a proactive issue, this is the rate at which action is taken. Provide the average time it takes from when a proactive opportunity is sufficient to warrant action and when appropriate action is taken. The rate will be considered to be the reciprocal of this average.

A note about Y_{if} versus Y_{ir}: Look at the number of times a degradation that could be detected from a proactive program are reacted to versus proacted against. The proportion of the time that a degradation problem is proacted against should be close to $Y_{ir}/(Y_{if}+ Y_{ir})$ if the time to address known proactive problems is roughly the same as the time to address reactive ones. Use this approximation for determining a target or for creating targets for future PNM programs or projects, or for estimating whether there is sufficient opportunity for a particular PNM solution in an area.

3.3. States

For the 11-state model that models the repair state of a section of plant, consider the following state definitions.

- w_n = working state at degradation level n. This is when there are no failures in the scope of the network being modeled, so service is working.
- Ip_n = indication of a proactive action. This is the state of the system in which an impairment has become severe enough that it is detected by your proactive system. As the model (in this formulation) assumes that the next action is to address the proactive issue, precluding a reactive one from happening meanwhile, the time to leave this state should be small, so it is likely best to assume the detection is severe enough to warrant action, too.
- Cp_n = customer impacted or triggering a threshold for proactive action. This is the state of the system in which the impairment that was detected proactively has now gotten severe enough that it will definitely be addressed now, either because a threshold was triggered, or a customer was impacted and indicated so.
- Rp_n = remote proactive fix attempt. This is the state in which a service center is working a proactive ticket.
- Dp_n = dispatched proactive fix attempt. This is the state in which the proactive ticket is generating a technician dispatch.
- Tp_n = technician working the proactive ticket. This is the state in which a dispatched technician is working the proactive ticket.
- Fr_n = failure, reactive state. This is the state under which a reactive service failure has occurred for the system being modeled.
- Cr_n = customer called in on a reactive failure. This is the state under which the customer has called in a problem, or an alarm has been indicated that requires repair.

- Rr_n = remote reactive fix attempt. This is the state of the system in which a call or service center is attempting to remotely fix the problem.
- Dr_n = dispatched reactive fix attempt. This is the state of the system in which a technician has been dispatched to fix the failure and service disruption.
- Tr_n = technician working the reactive ticket. This is the state under which the technician is actively working on the trouble ticket to fix the problem, after being dispatched.

3.4. Constants

Several constants need to be defined as well, because after a time in a state in the model, there may be choices for the next state, and a constant proportion of the exits from a state may enter a given state next, which must therefore be defined. For each proportion defined, there is a complement (one minus the proportion) that describes the probability that the next state is some complementary state.

- a = proportion of reactive tickets that are truck rolls or dispatched.
- b = proportion of remote repairs that close the ticket, not passing it off to a dispatch.
- c = proportion of repairs that don't resolve an impairment but close the ticket and complete a service repair (minimal repair, doesn't improve the plant).
- d = proportion of dispatch repairs that have to try again, either becoming a repeat, or perhaps don't have access to the location of the trouble.

Note we can add a sub-notation of p for proactive or r for reactive to each of these constants to have different proportions for reactive and proactive work. In cases where either p or r exist in the notation but not both, the assumption is that the notation without the qualifying subscript is represented by the factors missing the subscript, to simplify the equation noise.

There are additional constants that are not proportions but are important to the model.

- e = constant factor for the variable part of the degradation rate, a function of the state. This constant factor is not a proportion, and applies to the degradation model.
- n = working state number most closely linked, i.e., w_1, w_2, etc. This is the level of degradation in the plant: n=0 when the plant is brand new, pristine, and without impairments; as it ages, n increases to a high number representing a highly damaged and degraded plant which generates lots of impairments that need to be fixed.

4. Finding Parameters to Feed the Models

To obtain the parameters for the model, it is anticipated that operations measures can be translated to the model parameters suitably through mathematical transformation: Rates can be determined from frequency averages, and proportions can estimate probabilities. As such, these expected operations measurements, which should be available, can provide the basic information for the model. Additional parameters can be set as goals for programs to be evaluated or searched for tipping points.

Several operations measurements that are relevant to the model inputs include the following, though some cable operators may have different sources for estimates than those assumed here.

- Number of trouble calls per customer over a period of time.
- Number of system alarms per customer over a period of time.
- Number of proactive service impairments discovered per customer over a period of time.
- Number of other categories of sources of maintenance work per customer over a period of time, if any, and what makes them distinct from the categories above. We expect this is zero but need confirmation. If it is non-zero, we will need additional information about this category of maintenance work in parallel to the other categories, as below.

- For each of these above sources of work, the proportion of the work done to correct a shock failure versus a failure due to degradation.
- For all the sources of work except proactive, the proportion of these that were indicated by proactive systems but not addressed on time.

Several measures are for reactive work specifically.

- The average or expected time between when service is impacted to when a customer calls in to request help.
- The average or expected time between when service fails to when the alarm indicates there is a problem.
- The average or expected time between when the customer calls in to when someone in a call center or service center responds (include dispatch if it is possible to go right to dispatch or truck roll without touching the center).
- The average or expected time between when the alarm sounds to when someone in a call center or service center responds (include dispatch if it is possible to go right to dispatch or truck roll without touching the center).
- The average or expected time between when a call center or service center responds to a customer call to when it is resolved or decided to dispatch.
- The average or expected time between when a call center or service center responds to an alarm to when it is resolved or decided to dispatch.
- The average or expected time between when a customer call is dispatched to when repair begins.
- The average or expected time between when an alarm is dispatched to when repair begins.
- The average or expected time between when repair begins for a customer call event and when service is restored.
- The average or expected time between when repair begins for an alarm event and when service is restored.
- The proportion of customer calls that go to a call or service center versus right to dispatch (if any).
- The proportion of calls to a service center or call center that are then dispatched.
- The proportion of calls to a service center that are resolved in the center and do not repeat.
- The proportion of calls to a service center that are resolved in the center but end up generating a repeat problem.
- The proportion of dispatched repairs that generate a repeat problem.
- The proportion of dispatched repairs that do not generate a repeat problem.
- For reactive work that is dispatched and does not generate a repeat problem, the proportion of repairs that correct the source of the problem versus minimally repair the cable plant just enough to get service restored.

Likewise, some of these equivalent parameters are needed for proactive work too.

While the above parameters are likely available from operations management systems, we need to translate them into parameters for the model still. For averages, we convert those to rates. If there is information about the variance of these averages, then we can use that as further verification as to the reasonableness of the assumption of exponential distribution for the durations, or make adjustments accordingly. The proportions are directly usable in the model. However, some of the first few parameters on this list need translation as follows.

The first group of six estimates will tell us the parameters Y_{bs}, $n{*}Y_{if}$, and $n{*}Y_{ir}$. The fifth estimate(s) gives us a breakdown of the first four so that we can estimate Y_{bs}. Subtracting Y_{bs} from the sum of the rates for issues to address, we find the sum of the other two rates. Then the sixth estimate helps us estimate how much of this sum is attributed to Y_{if} and how much to Y_{ir}. The fact that these are averages only tells us an average value for the parameters multiplied by n. So, we rely on the estimate of the proportion of work that minimally repairs the cable plant versus improves it as a way to estimate the impact on plant impairments. That helps us estimate the value of n=1. From the proportion of work that is based on shocks versus degradations, we can estimate the number of impairments in the plant utilizing methods from software reliability. Then, we can estimate the Y_{if} and Y_{ir} from our estimate of n and the sum of the rates we estimated above.

The age of the plant being modeled is somewhat important as we estimate the time to an impairment entering the plant and being discoverable. The distinction of those two steps is not likely important to separate, but we do want reasonable estimates for our model.

Model Use and Use Cases

Many use cases can be addressed with just a steady-state analysis of the 11-state repair model by comparing two or more sets of parameters. Very quickly, the future model state no longer depends on the starting condition, so a steady-state result is sufficient, and easy to solve using traditional methods. The overall service or plant availability can be determined by examining the steady-state probabilities of the model. A current state and future state model pair can be compared, and more complicated combinations can be calculated too.

Decisions around changes to operations are usually assessed by cost-benefit analysis, so a cost model is very important. Costs often have a fixed component, and a variable component that is often a function of time. So, we need to know the time spent in each model state, and the number of visits over time to each state, as well as the costs involved.

To keep the analysis simple, we will consider the average costs of visiting each state, so that the cost includes the fixed and variable costs together.

In the models we encoded and ran, we broke out the costs to the visits to each of the states so that the cost of repeating cycles can be handled in detail. The basic set of cost parameters we use are as follows.

- 10 = cost of goodwill from a reactive failure
- 20 = cost of goodwill when a customer calls to complain
- 30 = cost of a remote fix
- 50 = cost of a truck roll fix
- 50 = cost of the actual repair, separate from truck roll
- 10 = cost of a proactive failure, likely system cost
- 5 = cost of a proactive alarm, likely system cost
- 20 = cost of a remote fix to a proactive problem
- 30 = cost of a dispatched technician fix to a proactive problem
- 50 = cost of the actual proactive repair, separate from the truck roll
- 0 = cost of being in a fully working state, likely 0

Note we could add a negative cost to the fully working state if we wanted to reflect a benefit from being in that state.

For the models we run later, we asked for input from operators, and a few shared some numbers we were able to use. We selected baseline transition rates and constants to reflect what we found from the operators who responded.

The state transition rates we used are as follows.

- Y_{imp} = 20/(365*24) = impairment arrival rate (ignored in 11-state)
- Y_{bs} = 200/(365*24) = base failure rate for shocks
- Y_{if} = 50/(365*24) = impairment failure rate, multiply by n for degradation state
- Y_{cc} = 4/24 = customer call rate
- Y_{rr} = 4 = response to reactive call rate
- M_t = 1 = dispatch rate
- M_d = 1/2 = dispatch repair rate
- M_r = 4 = remote repair rate
- Y_{ir} = 2/365 = impairment reveal rate, multiply by n for degradation state
- Y_{pc} = 1/24 = proactive alarm rate
- Y_{pr} = 4 = response to proactive rate

The baseline constants we use are as follows.

- a = 0.0 = proportion of reactive tickets that are immediate truck rolls
- b = 0.75 = proportion of remote repairs that close the ticket
- c = 0.25 = proportion of dispatch repairs that don't resolve an impairment but close the ticket
- d = 0.17 = proportion of dispatch repairs that have to try again
- e = 0.1 = acceleration factor for impairment arrival rate for third order effect
- ap = 0.0 = proportion of proactive tickets that are immediate truck rolls
- bp = 0.75 = proportion of proactive remote repairs that close the ticket
- cp = 0.25 = proportion of proactive dispatch repairs that don't resolve an impairment but close the ticket
- dp = 0.17 = proportion of proactive dispatch repairs that have to try again

1. Aligning the Repair Model to the Degradation Model

A simple approach for using these models for a long-term cost analysis is as follows. See the top and bottom halves of the 11-state model in Figure 1. The top and bottom lobes (in green and brown) each represent a repair cycle. The top lobe (green) is for proactive work, and the bottom (brown) is for reactive.

When starting in the working state, the probability of entering the proactive lobe next (before entering the reactive lobe) is $(n^* Y_{ir}) / (n^* Y_{ir} + Y_{bs} + n^* Y_{if})$ so the complement probability, of entering the reactive node next, is $(Y_{bs} + n^* Y_{if}) / (n^* Y_{ir} + Y_{bs} + n^* Y_{if})$. Each lobe though can cycle within, and follow different paths in each lobe, so obtaining the costs requires more work.

For the 11-state repair model, we can calculate the expected time in each state from the transition rates, and solve the model for steady-state probabilities to get the percent of time in each state. By taking the probability of being in a state, dividing by the expected time in that state when visiting it, and multiplying by the cost of being in each state, we can estimate the expected cost of a cycle of the model. Dividing by the expected time for a cycle, we get the cost per unit of time of being in that 11-state model. Because we can do this for all degradation states, and thus the 11-state model that aligns with each degradation state, we can find the cost per unit time as the degradation proceeds.

Let Q be the transition rate matrix, where $Q_{i,j}$ = transition rate from state i to state j when i≠j, and the negative of the row sum of the off-diagonal values otherwise for i=j so that the row sums to 0. This is the standard transition rate matrix for stochastic models.

PIa is the steady-state probability results from the Q matrix, representing the probability of being in each state i over the long run, meaning regardless of the starting state.

$-Q_{i,i}$ is then the transition rate out of a state i. Thus $-1/Q_{i,i}$ is the expected time in a visit to state i, or the sojourn time in state i.

Thus, we can get $1/(-Q_{i,i} * PIa_i)$ to be the circle back time, or the time between visits to a state i. Also, this is 1 over the number of visits per unit of time. Therefore, -Qi,i * PIai * Ci is the cost per visit to a state multiplied by the visits per unit of time, so is the cost per unit of time for each state, which we can sum to get an estimate of the cost per unit of time for the whole model.

This lets us calculate the 11-state model for every level of degradation we want to model, then apply a disjoint model of degradation level to the cost estimates to get an overall cost of the plant. Further, we can separately model degradation to model how the cost changes over time, and then plan for optimal maintenance of this cost profile. Likewise, this allows us to use a cost per unit time as a way of comparing different model settings which may represent two use cases to compare, as we will show later.

2. Modeling Plant Degradation as it Relates to Repair

The 11-state repair model can be connected to a degradation model as we describe here.

First, we model the degradation level of the section of plant as a birth-death model, meaning that the state increases or decreases by 1 at each state transition. If we let the first state represent the pristine cable plant, then a birth represents a degradation, and a death represents removing the degradation from the plant to return it back to good as new for the part of the network that was affected.

To use a degradation birth-death process model, we need to determine the degradation rate (birth) and improvement (death) rate based on the state n. Assume the degradation birth-death model has 100 states numbered n=1 to 100. This assumption is arbitrary and can be easily relaxed; but for our purposes, 100 states should be enough to model reasonable levels of cable plant degradation.

For degradation, we define a constant rate of degradation Y_{imp} added to a rate that depends on the state n, $e*(n-1)*Y_{imp}$. Note we use n-1 instead of n because the first state 1 will then have a degradation equal to the base constant rate. Also, note that this degradation rate applies to states n=1 to 99 as state 100 can’t degrade any more.

For the improvement rate, we have to determine the probability of being in a repair state that allows improvement, times the rate at which an improvement happens. There are two such states in the 11-state model. This rate, dependent on the current state n, is then

$$Pla_{tr,n}*(n*Y_{if}/(n*Y_{if}+Y_{bs}))*Md_r*(1-d)*(1-c) + Pla_{tp,n}*Md_p*(1-dp)*(1-cp)$$

Note the improvement rate applies for n=2 to 100. Also note we allow M_d to be different for proactive and reactive in this equation, so adjusted the notation respectively to Md_p and Md_r.

3. Reactive and Proactive Relationship with Failure Rate

The impairment failure rate and impairment reveal rate are related in a potentially complicated way, depending mostly on the operations data feeding the model. Some of the impairments will lead to failures, and some will be revealed and acted on before becoming a failure due to a proactive capability or program. This relationship requires a reasonable model to trade off one for another when comparing reactive only to proactive and reactive programs mixed.

First, we assume that the failure rate is a combination of multiple failure causes each occurring with its own rate, and the net failure rate is approximated by the sum of rates. Incidentally, this summation simplicity is one reason we prefer to work with rates rather than time between failures; one is the reciprocal of the other. Due to this assumption, we can discuss one part of this failure rate for a single failure cause. Note that a failure cause in this case can specify a network component type and failure mode, or a group of either, but is specific to the scope of the proactive maintenance that is possible.

Say the rate at which a given problem appears in the network being modeled is Y_1. For simplicity, this rate can include the n factor for degradation within it. If the proportion of these problems that get caught proactively is p_p and the reactive proportion is p_r such that the two sum to 1, then a simple way to split them is to let the proactive part be p_p*Y_1 and the reactive portion being the complement. As long as the proportion is applied randomly, the results are two thinned Poisson processes which are each Poisson processes. They also can add back together, which makes the modeling simple to manage and adjust.

However, it is far more likely that we have the all-reactive operational data and are looking to examine the impact of changing some reactive work to proactive. Doing this requires we reverse the process we just described earlier by reducing the reactive rates by the amount that converts to proactive (thinned Poisson process), and then pushing the proactive work to the proactive part of the 11-state model. If we believe that the shift of work would change the other reactive rates in the model, then adjustments to those are in order. However, for our scenarios in this paper, and likely for many operations cases, the secondary effects are likely negligible. Therefore, for comparing reactive only to proactive-reactive mix environments, we can simply divert some of the reactive work to the proactive side, and reduce the rates as thinned Poisson processes so that the all reactive case has a rate Y_{if} that is equal to the proactive-reactive rates $Y_{if} + Y_{ir}$ which is a simple comparison.

We can write this relationship as $Y_{if,r} = Y_{if,p} + Y_{ir,p}$ adding the r for the reactive-only model parameter, and the p for proactive-reactive model parameters. We will apply this method in the models in this paper unless otherwise stated.

Model Study Results

We will now use these models to study several use cases for reasonable parameter estimates obtained from discussions with operators willing to share their knowledge.

- Advantages of PNM – We will run a sensitivity analysis on the conversion of proactive and reactive to reactive only in the 11-state model and determine if PNM is a strong advantage.
- Reducing Repeat Rates – We will start with a simple repeat rate model where we model a high repeat rate and then a much lower rate but with different repair times and costs, all using the 11-state model.
- Degradation Over Time and Optimizing Cable Plant Rehabilitation – We will model degradation over time according to the degradation model outlined before, but with different settings for degradation and removing the degradation causes encountered in repair. The analysis will reveal

the behavior over time of the degradation model, which we can consider against best practices for maintenance.

1. Reactive-Only versus Proactive and Reactive Mixed

In our first use case study using the 11-state repair model, we compare the baseline parameters for a mix of proactive and reactive, then adjust the proactive work to be reactive as described earlier but allowing for some portion of the proactive work to turn into reactive work, from 50% to 100% in increments of 5%. This allows for situations where the conditions generating proactive work may not always be conditions that turn into reactive work. As little as 50% is considered, but as much as all of it is in the extreme.

We allowed the baseline rates and constants to vary so that we could generate Monte Carlo runs of different parameters for comparison. Rates were allowed to vary from 1/3 to 3 times the base rates. The constants were varied by 1/3 of the proportion less than 0.5, on either side. For example, if a constant was 0.75, then it was allowed to vary by 1/3*(1-0.75) = 0.25/3, so can be 0.75 ± 0.25/3. We then generated uniform random numbers in the range that each was allowed to change and created 100 different sets of settings. We then applied each model condition to the 100 setting sets to generate 100 model results for each of the 11 conditions. We then used these results for comparisons as follows.

See Figure 7 for the availability results of the 11 different model results. We generated box plots of the 100 Monte Carlo availability results for each model condition, assuming the starting state was in degradation state n=38, as that was a state close to what one operator described as their average condition. See that almost all of the availability results for a proactive and reactive repair system were better than any of the results for the reactive-only results, even when the reactive fixes of proactive problems were half what the proactive work would have been. In other words, from a service availability standpoint, proactive repair is much better in general.

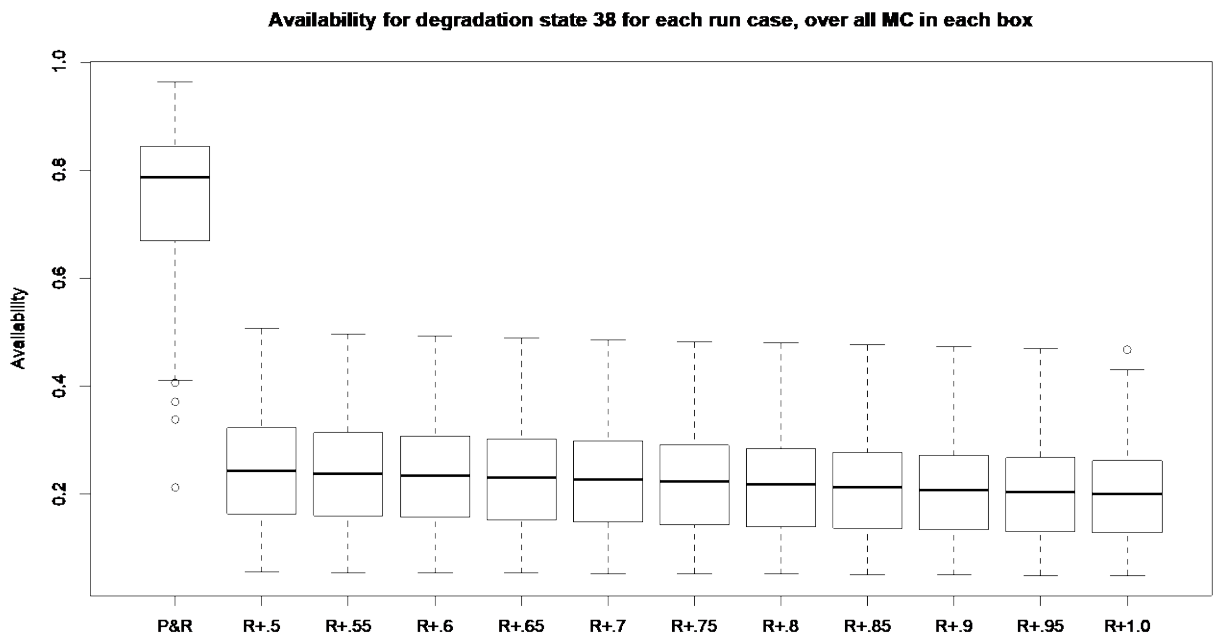

Figure 7 – Service availability over 100 Monte Carlo runs for a proactive and reactive maintenance system, and reactive-only where the proactive work generates 50% more to 100% more reactive work when not handled proactively.

Now examine the cost rate results in Figure 8. The data in these box plots are cost rates from the same runs as in Figure 7, so consider a degradation rate at 38. Note that the repair system that has proactive and reactive work has a much lower cost rate as indicated by the box plots, compared to the various reactive-only repair systems. Once again, we have some evidence that in general a mix of proactive and reactive repair will cost less than a reactive-only repair system.

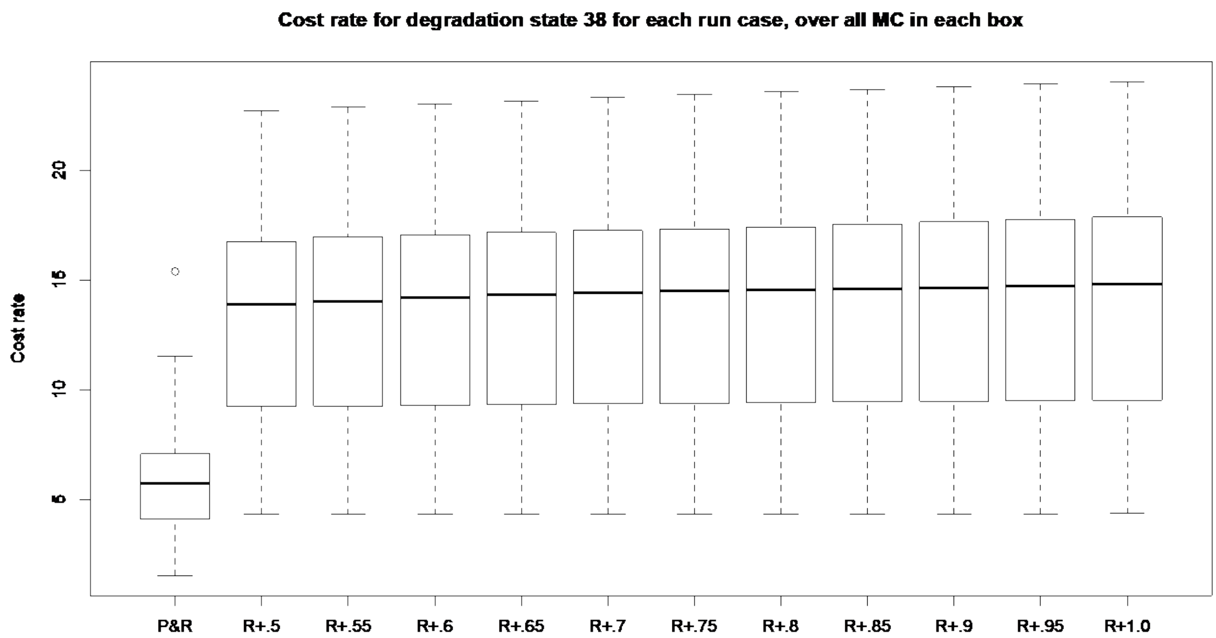

Figure 8 – Cost rate over 100 Monte Carlo runs for a proactive and reactive maintenance system, and reactive-only where the proactive work generates 50% more to 100% more reactive work when not handled proactively.

To check whether the degradation level of 38 is somehow special or not, we provide box plots of all degradation levels, comparing the best case reactive-only repair to the mixed proactive and reactive repair solutions, in Figure 9. Because we take every model run and calculate the difference in the cost rates of the two model settings, we can say that any negative values would be due to a reactive only solution being better than the proactive and reactive mixed system. But look at Figure 9 and see that there are no box plots with even extremes in negative cost differences. This indicates that, at least for the settings run in the Monte Carlo model runs, having a PNM element to a repair process is cheaper than not having one.

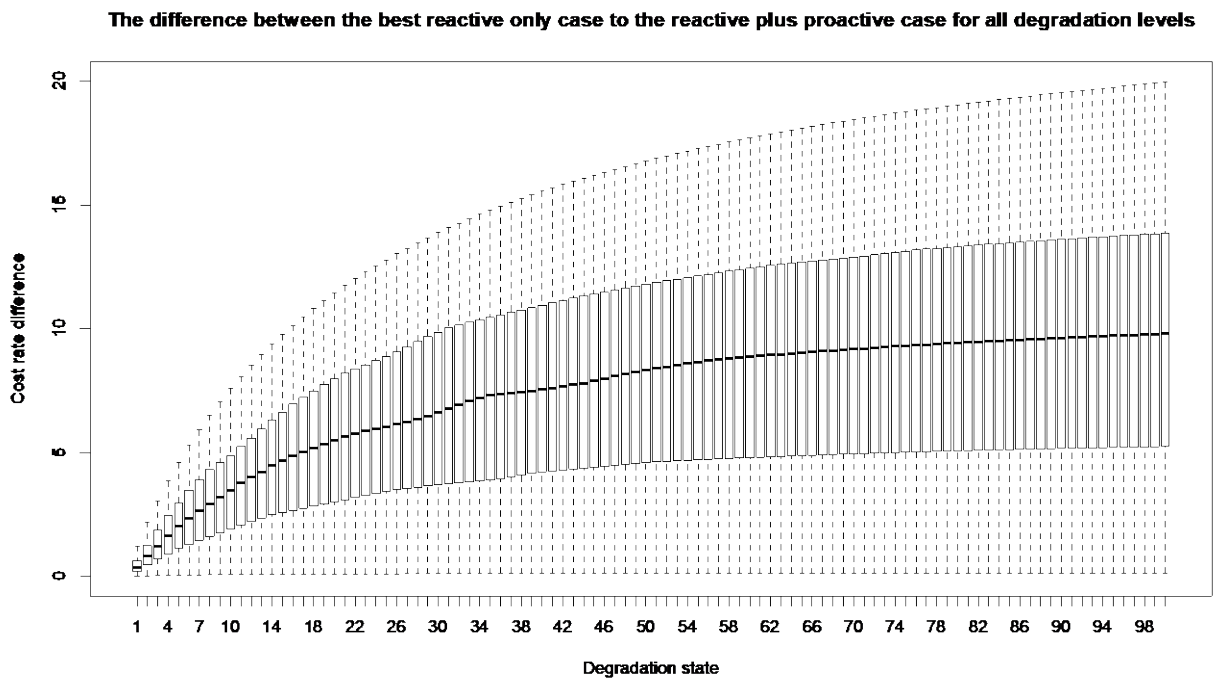

Figure 9 – Cost rate difference between the best reactive case to the proactive and reactive case for add degradation states, over 100 Monte Carlo runs.

2. Reducing Repeat Rates

One operator expressed a concern about their high repeat rates, so we decided to demonstrate how to use these models to assess whether it makes sense to let technicians take longer to troubleshoot and repair problems so as to reduce repeat rates, or not. The baseline repeat rate we used was 0.17 (3xR), and we examined competing cases where the repeat rate was 0.17/3, but at a cost of increasing the time and cost of repairs by 10% (LR1) to as much as 110% (LR11). Once again, we generated 100 Monte Carlo runs of the model, letting the rates and constants vary as before, except for the repeat constant.

See Figure 10 which has box plots of the results for the high repeat rate first, then the lower repeat rates with longer repair times as indicated, all for degradation level n=38 as before. As might be expected, the impact to service availability is very minimal, and decreases as repair times take longer. However, a good repair may mitigate this impact by restoring service quickly, then spending more time to find and remove the root cause of the problem.

Now look at Figure 11, which shows the costs per unit of time for the same runs. The high repeat rate has the cost shown in the first box plot, but notice that the first few lower repeat rate cases have lower costs per unit of time, suggesting that spending up to 30% or so extra time (cost) on the repair will still result in lower cost overall. More time spent in the field correcting problems will pay off in the long run.

To show that the degradation state of n=38 is not so special, we calculated the difference in cost per unit time between the high repeat rate case and the low repeat rate but spending an additional 110% of time troubleshooting. While the latter mostly has a much higher cost per unit of time, it isn't always the case. There may be times where it is still cheaper to spend twice the time to find and fix troubles.

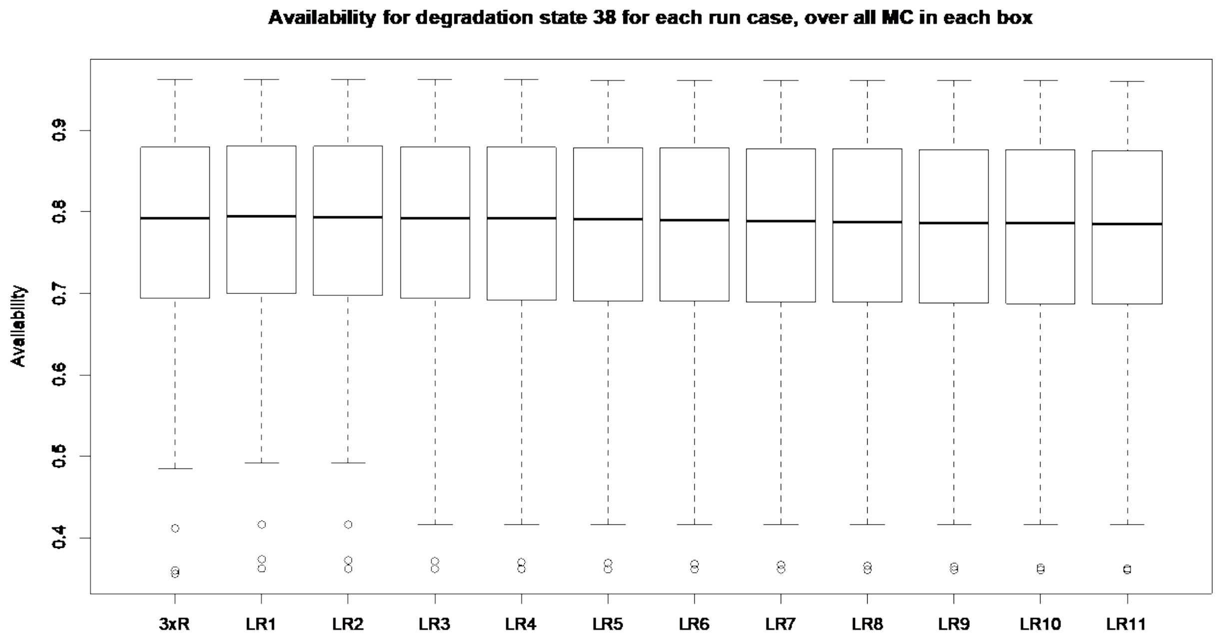

Figure 10 – Service availability over 100 Monte Carlo runs for a proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost to 110% more time and cost.

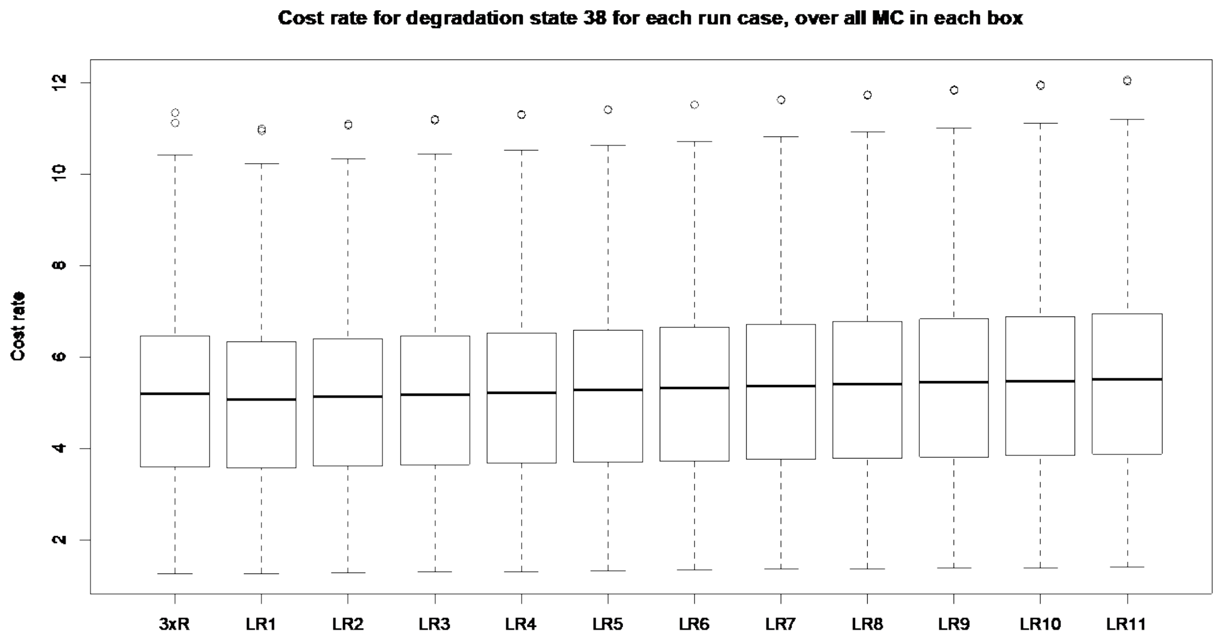

Figure 11 – Cost rate over 100 Monte Carlo runs for a proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost to 110% more time and cost.

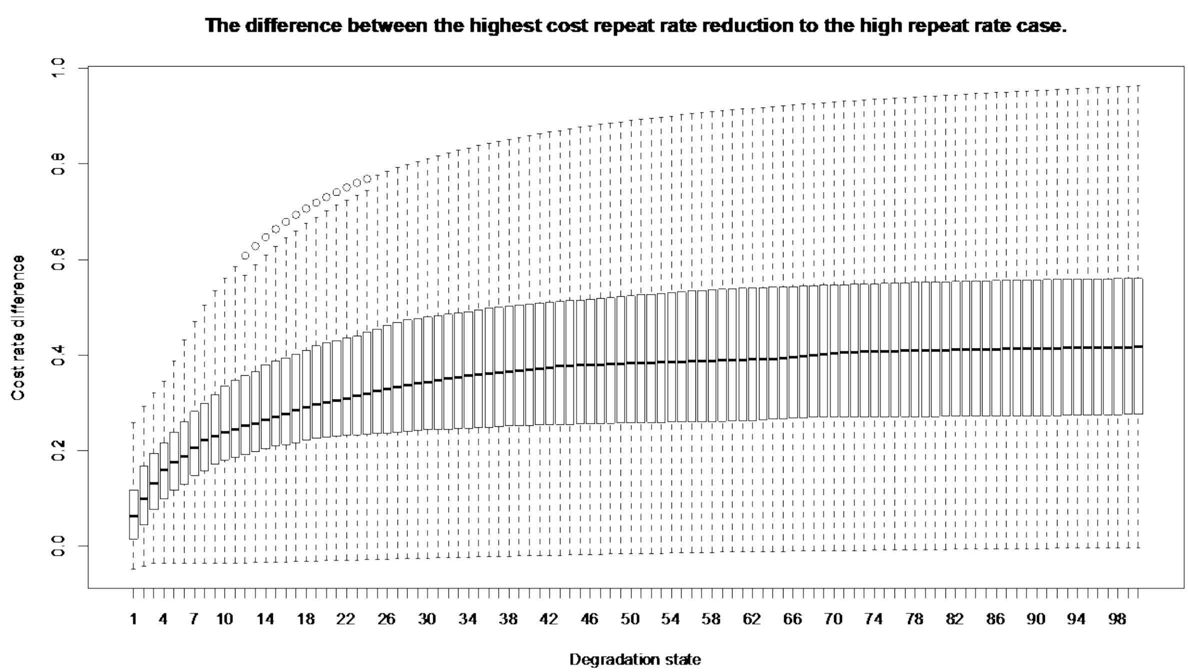

Figure 12 – Cost rate difference between the proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 110% more time and cost, over 100 Monte Carlo runs.

Last, we took the cost difference between the high repeat rate case and the low repeat rate case that requires spending only 10% more time and cost troubleshooting. That result is in Figure 13. Notice that the high repeat rate case is always more expensive, and in some cases significantly so.

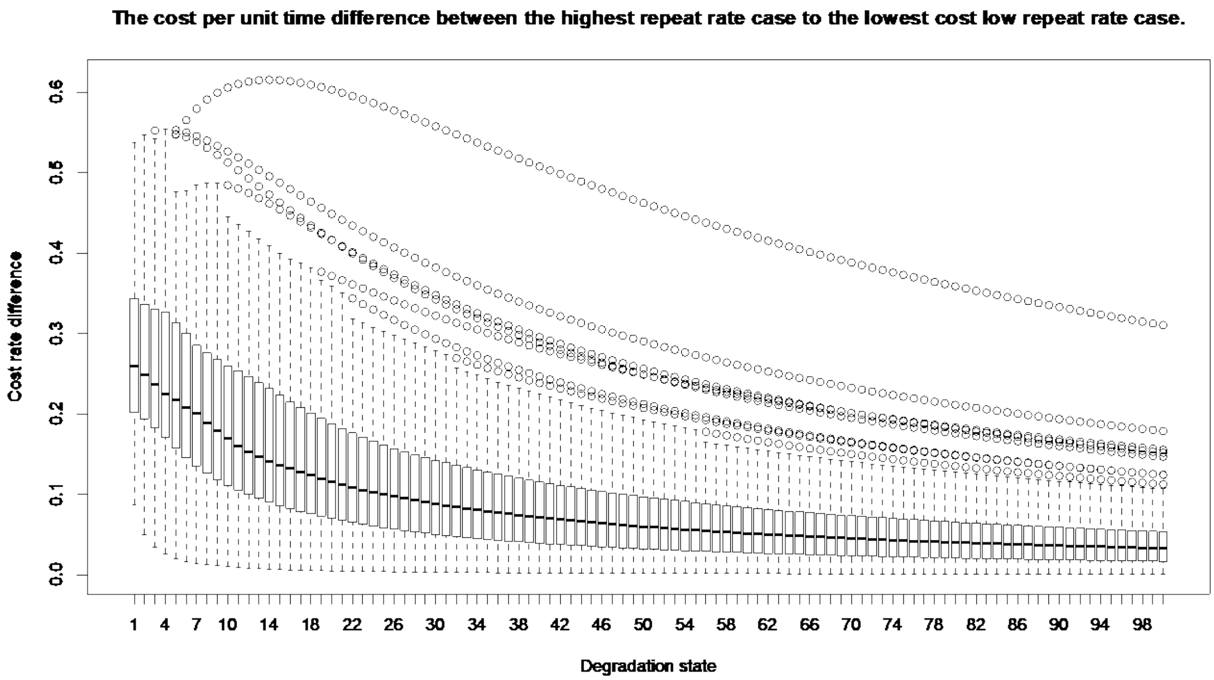

Figure 13 – Cost rate difference between the proactive and reactive maintenance system with 17% repeat rate, and 1/3 lower repeat rates obtained by spending 10% more time and cost, over 100 Monte Carlo runs.

3. Degradation Over Time and Optimizing Cable Plant Rehabilitation

We can model the probability distribution of a future state or estimate the past costs from an assumed start condition and age and parameters of current cable plant to estimate the lifetime cost functions for a section of plant. There are multiple uses for these results, but a major one is to use these results to set the optimal replacement times for the cable plant [2]. It is known that for systems that degrade, when the total cost of the system divided by its age (so the cost per unit of time of the system) is lower than (overtaken by) the cost per unit of time of the repair, then the overall cost per unit of time of the system begins to increase, so it is the optimal time to replace or fully repair it.

We run the 11-state repair model over all 100 degradation levels, then plot the availability and cost rates for each in Figure 14, and Figure 15 respectively. Note as should be expected that the availability decreases with degradation, and cost increases.

But as we start in a specific degradation state but let the birth-death model transition over time, we get probability distributions for the degradation state in the future. We show the cost per unit time, then the probability distribution functions for the degradation levels, starting in degradation state n=38, for all four combinations of settings of 20 or 10 degradations on average per year, and 25% or 12.5% of repairs not fixing the root cause of the problem (not to be confused with repeat rate, but correcting a degradation 75% or 87.5% of the time), at an age of 60 months, in Figure 16 through Figure 23.

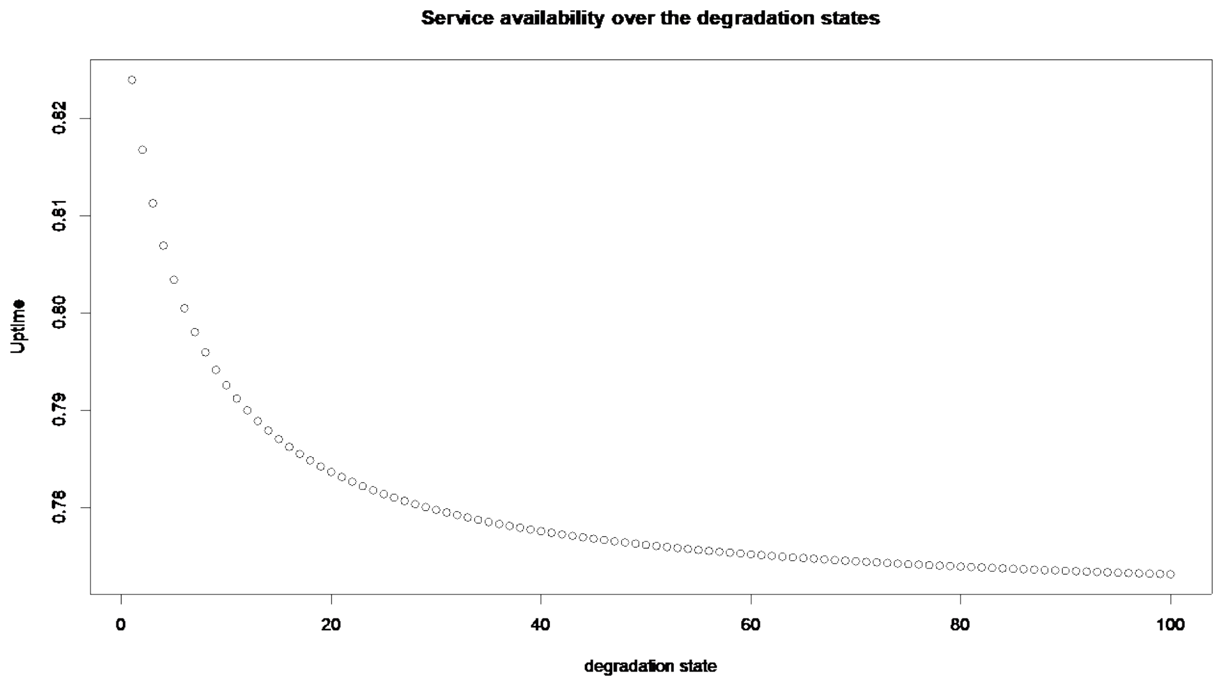

Figure 14 – Service availability over the degradation states, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.

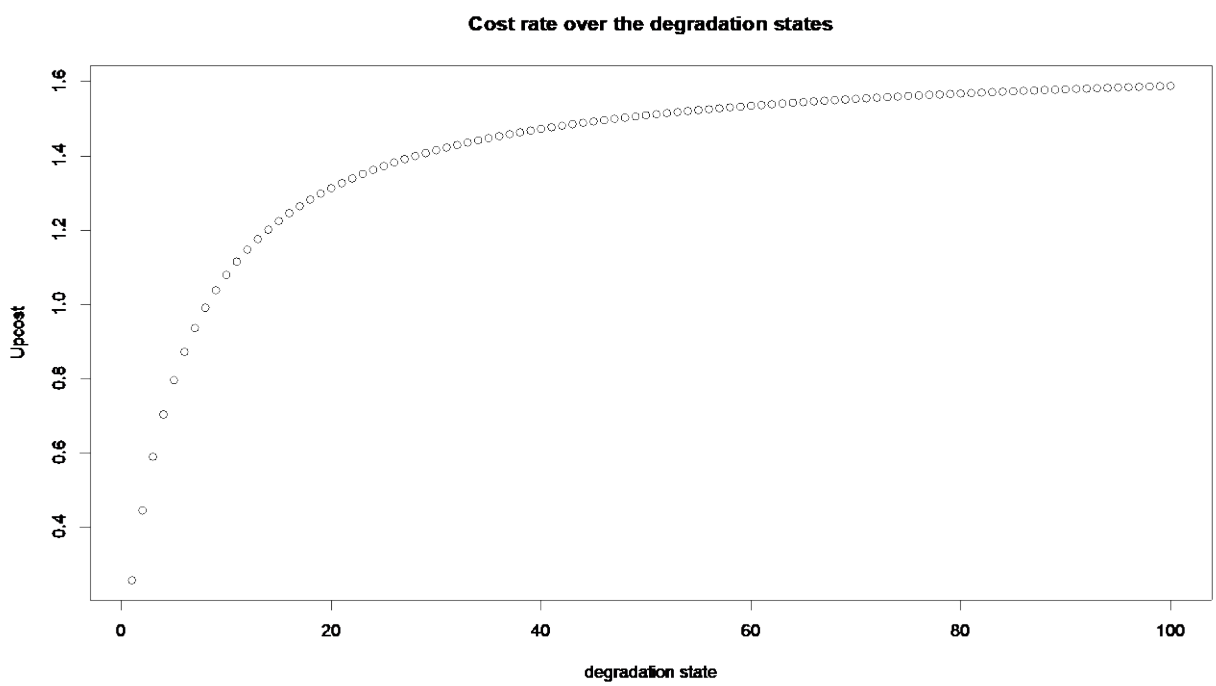

Figure 15 – Maintenance cost rate over the degradation states, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.

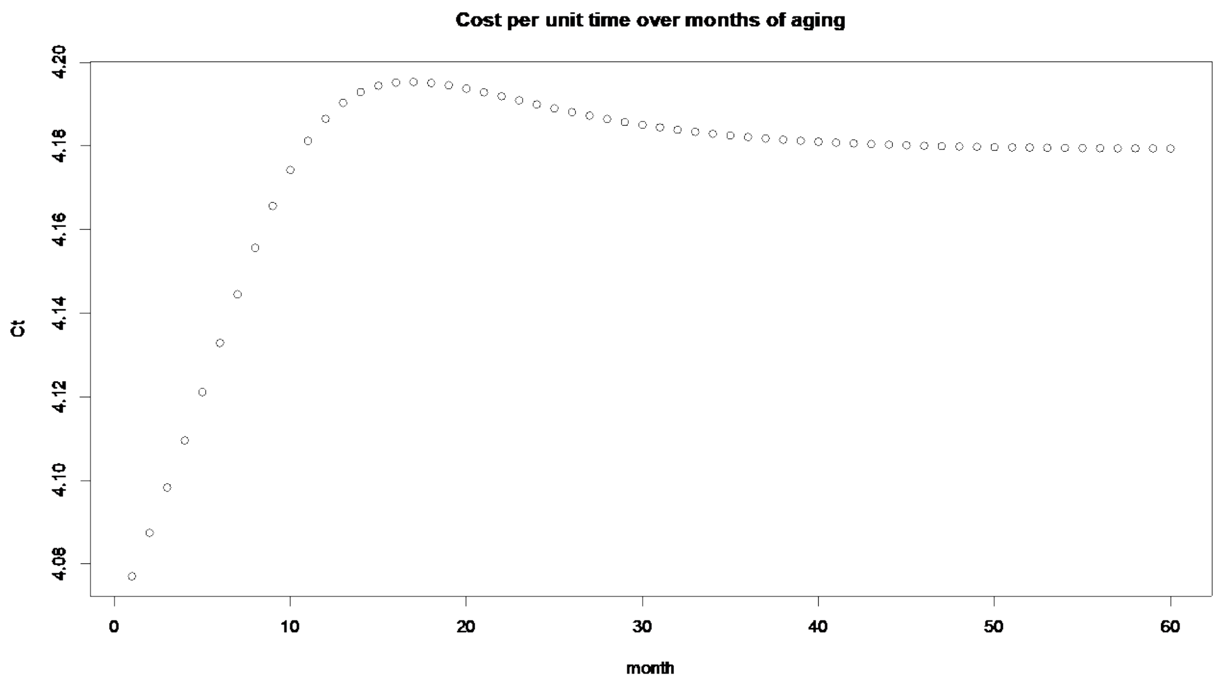

Figure 16 – Cost over time per month starting in degradation state 38, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.

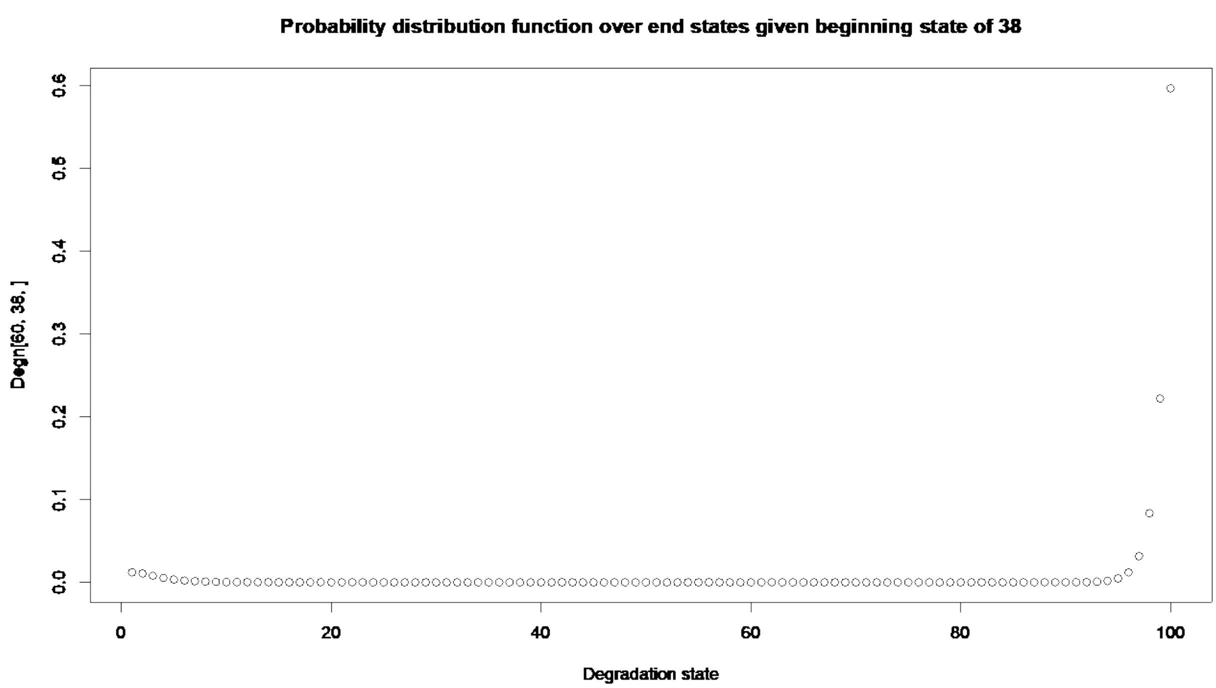

Figure 17 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.

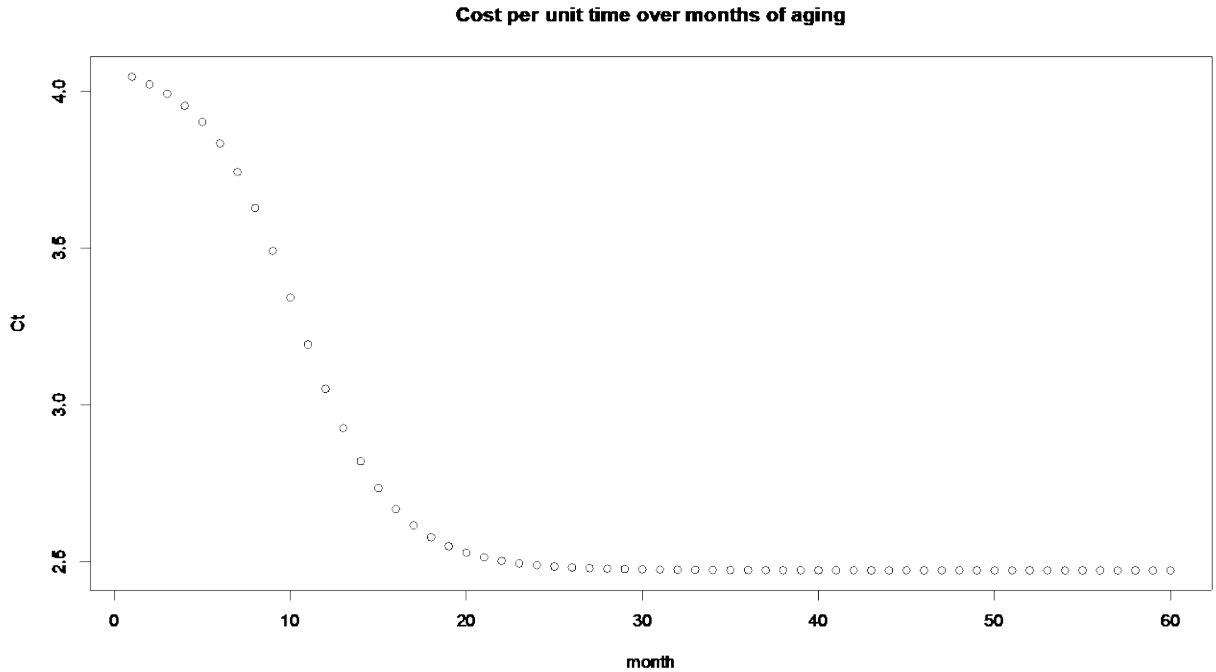

Figure 18 – Cost over time per month starting in degradation state 38, with constant degradation of 10 per year, and 25% of repairs not fixing the problem.

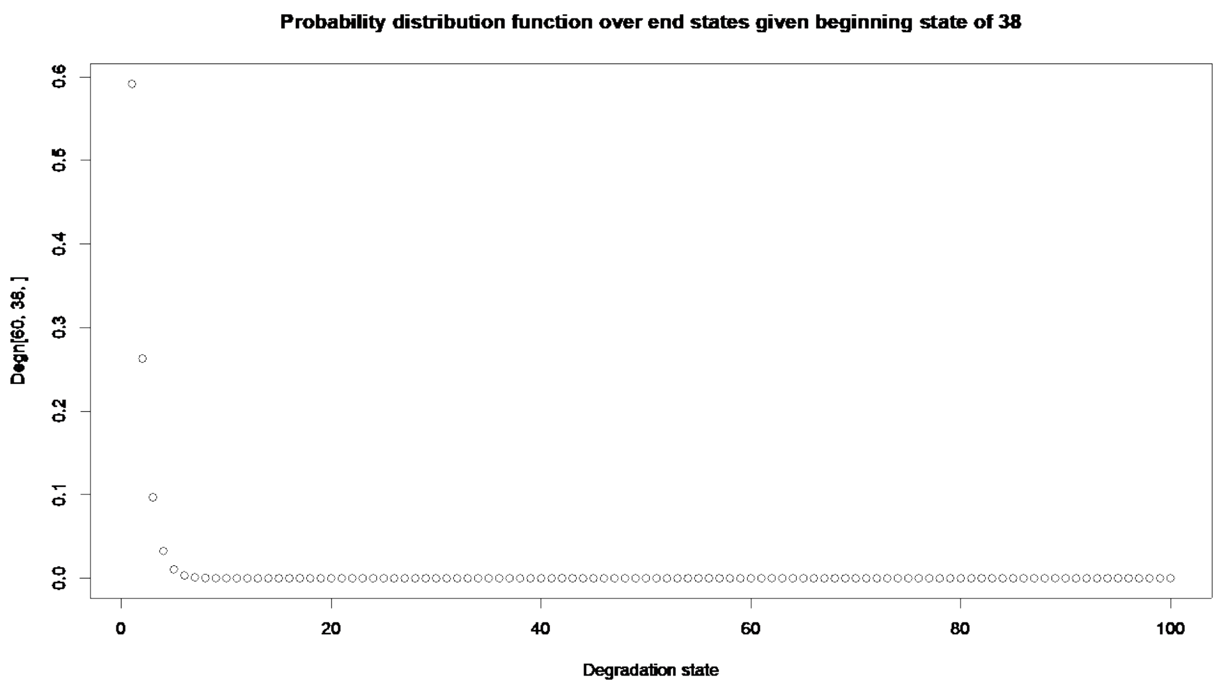

Figure 19 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 10 per year, and 25% of repairs not fixing the problem.

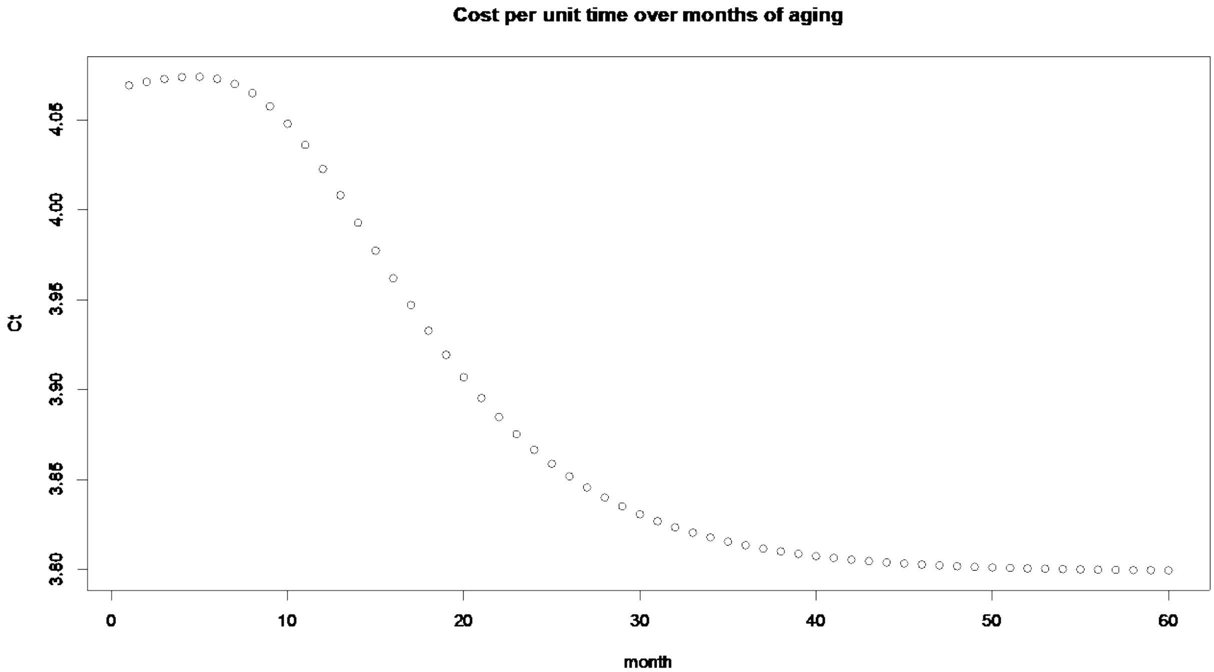

Figure 20 – Cost over time per month starting in degradation state 38, with constant degradation of 20 per year, and 12.5% of repairs not fixing the problem.

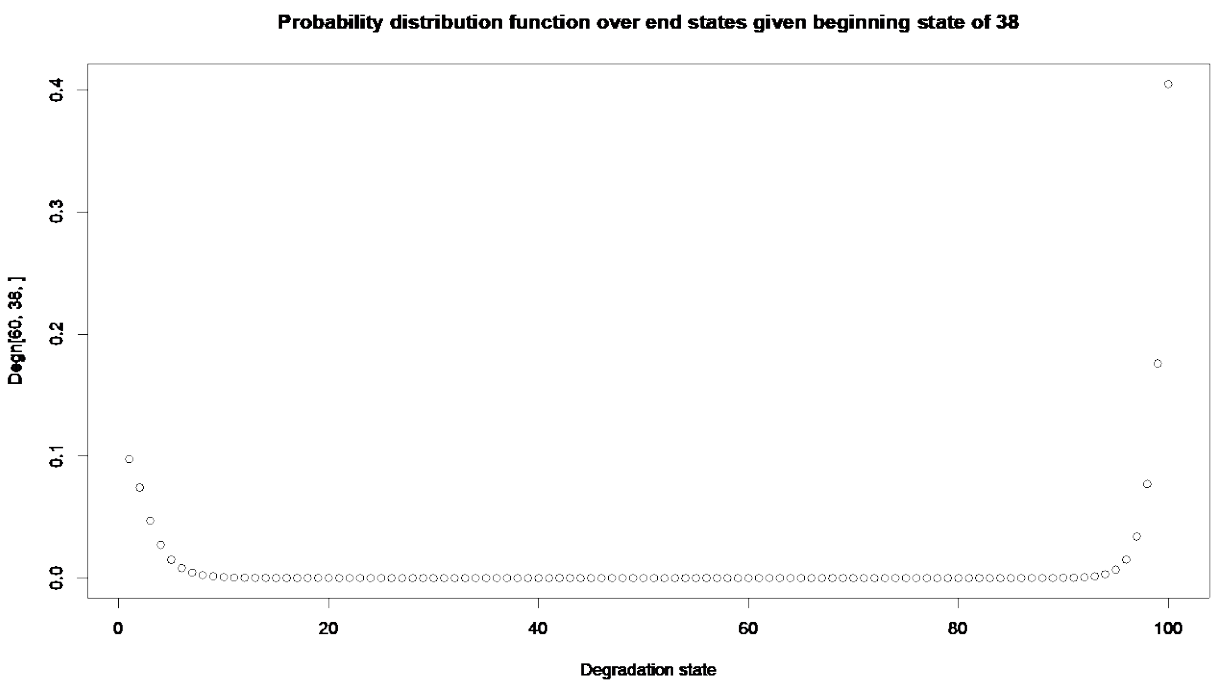

Figure 21 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 20 per year, and 12.5% of repairs not fixing the problem.

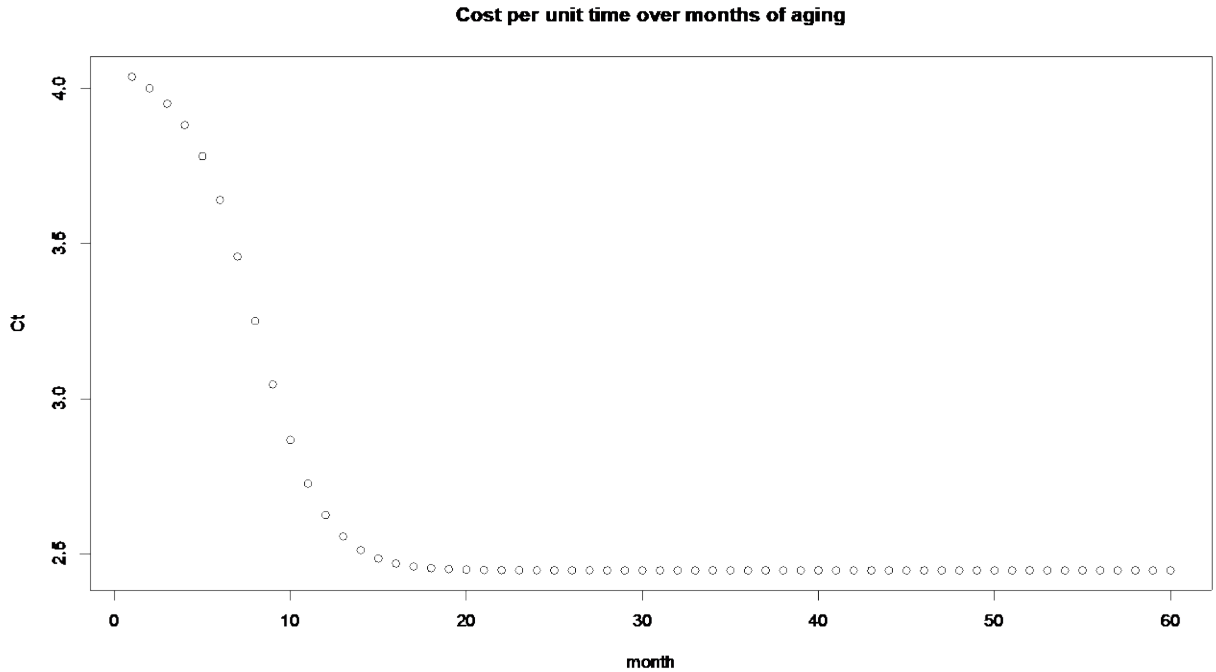

Figure 22 – Cost over time per month starting in degradation state 38, with constant degradation of 10 per year, and 12.5% of repairs not fixing the problem.

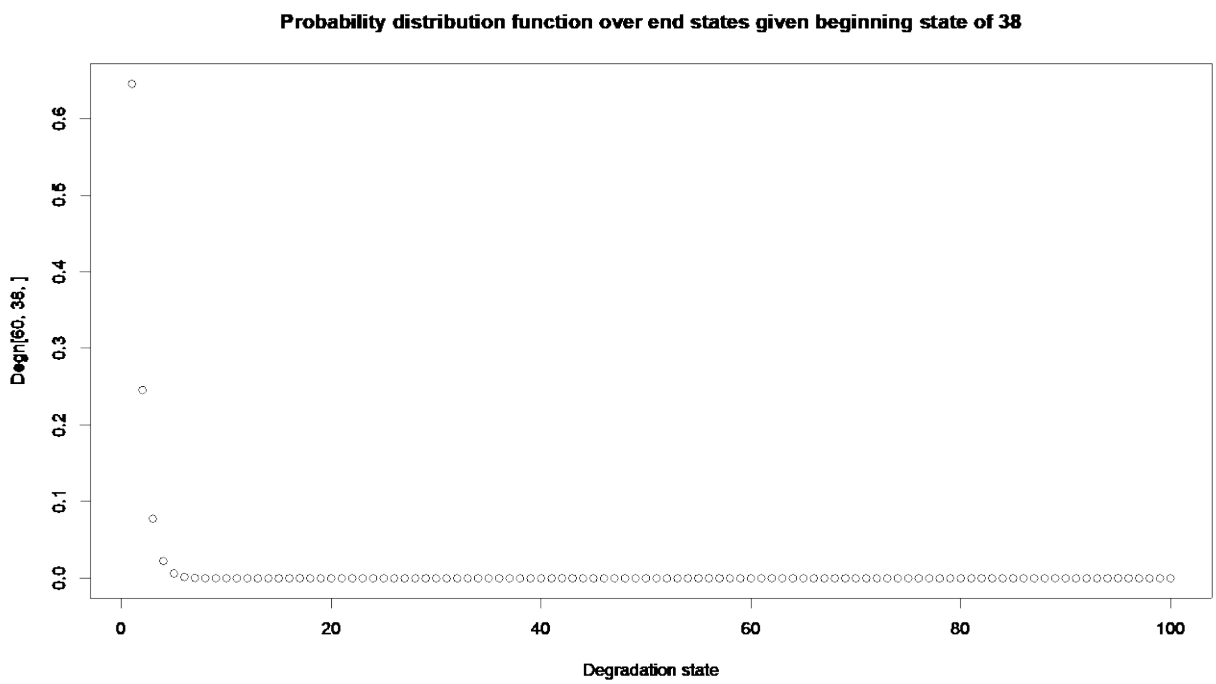

Figure 23 – Probability distribution function at 60 months, starting in state 38, with constant degradation of 10 per year, and 12.5% of repairs not fixing the problem.

Examining the probability distribution functions, we see a system that gravitates toward a low degradation level or a high degradation level, depending on the parameters. The system does not stabilize to a steady-state distribution of degradation states, but rather is likely to move toward the extremes. But there are cases where it could be at either extreme under the same settings. See Figure 21 for an example of such settings, where a low degradation rate is highly likely, as is a high degradation rate, but not so for intermediate states. Also, the costs per unit of time seem to stabilize. But this is likely due to the system stabilizing at an extreme. If that extreme is a high degradation level, then it could be that the degradation is stopped only by the finite states of the model, and that a real system would continue to degrade to alarming levels.

However, this movement toward a high degradation level can be avoided by PNM practices. Another well-known repair optimization result for systems like this is that it is almost always better to do condition-based maintenance over time-based maintenance. If we can examine that the rate of trouble tickets is increasing, and many of them are from degradation causes, then it might be time for a plant sweep or a plant rehabilitation (even if not proven to be the optimal time for it).

See Figure 24, and Figure 25 for the probability distribution function of states at 12 months, starting in degradation state 1, and state 20, respectively. These plots are both under the worst case of high degradation rate and lowest chance of correcting the degradations at repair opportunities. Yet after 12 months the systems are still likely in low degradation levels, because they started in low degradation levels. This suggests that removing degradations from the plant, and keeping them low, can keep costs low. An annual sweep might be sufficient, but monitoring for conditions and triggering a cleanup effort when degradations are appearing might work even better.

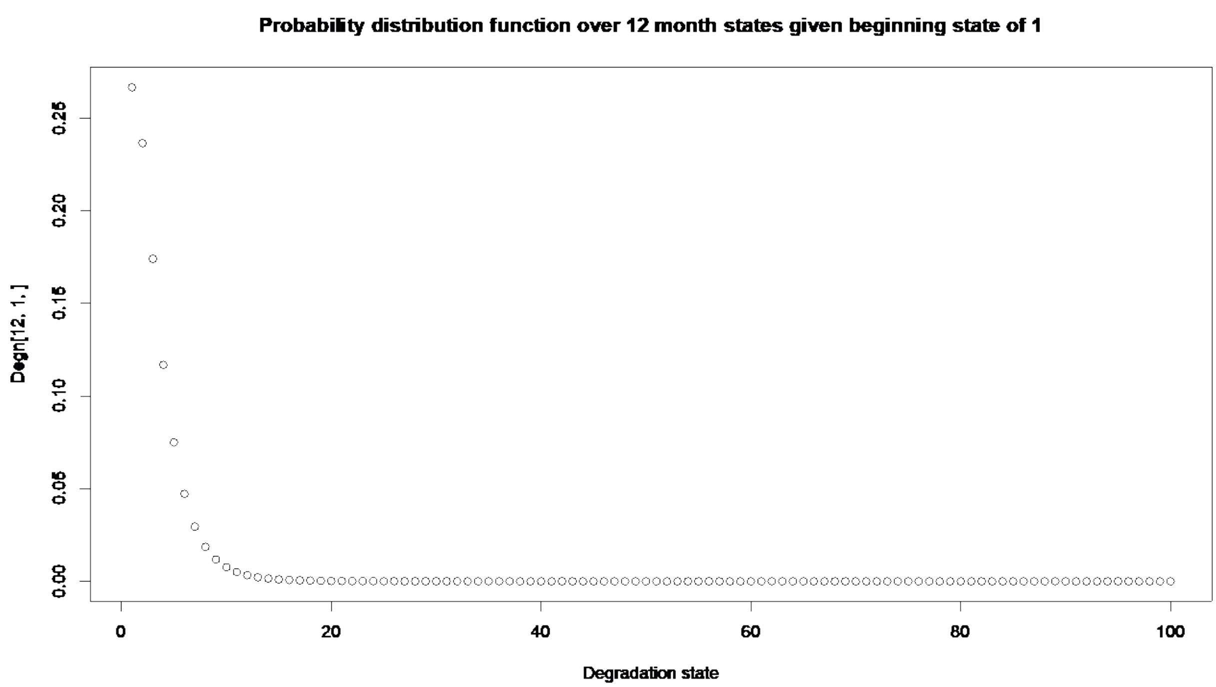

Figure 24 – Probability distribution function at 12 months, starting in state 1, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.

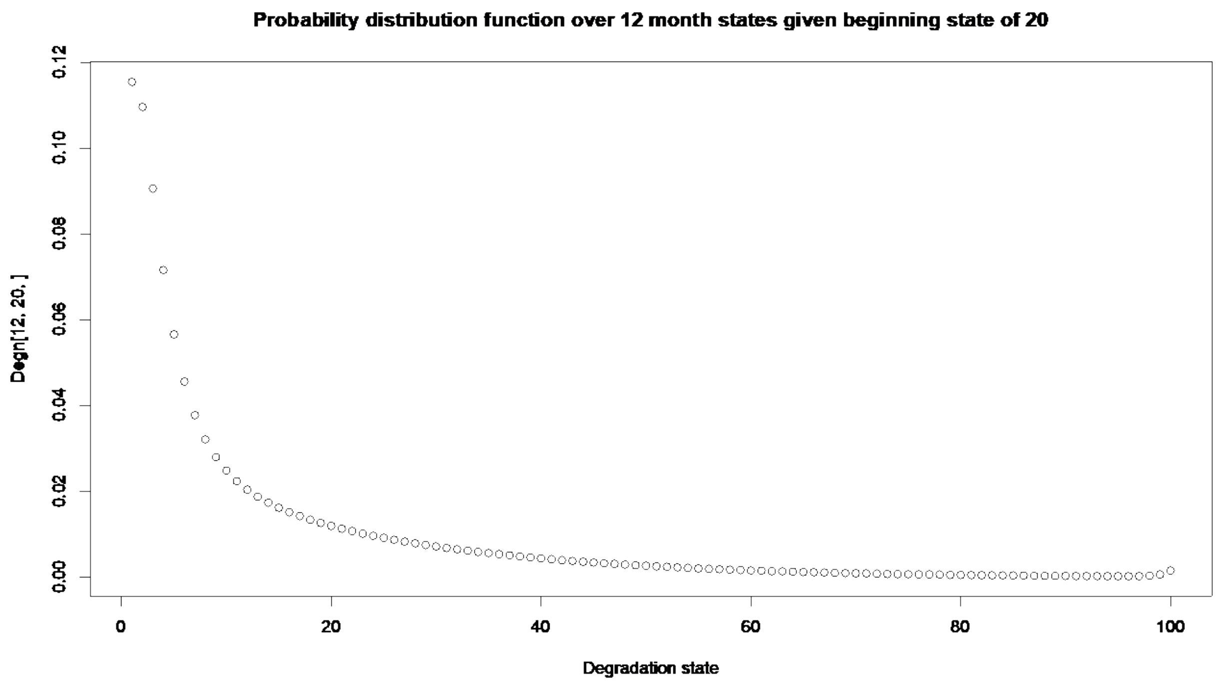

Figure 25 – Probability distribution function at 12 months, starting in state 20, with constant degradation of 20 per year, and 25% of repairs not fixing the problem.

Conclusions

While the models presented in this paper are general examples intended to be tuned to specific operator needs, and the analyses are examples to consider for applications, the results are still of use in a general sense.

1. Generalizing Results

These general models suggest that proactive maintenance can reduce maintenance cost overall and keep service availability high. Further, it can be cost effective to spend more time and technician cost correcting problems to avoid repeat troubles. Also, there is evidence that cable system repair can shift from maintaining cable plant at high quality to having accelerating degradation, but that monitoring or annual sweeps can keep operations costs lower.

2. Operator Uses

In this paper, we demonstrated several use cases for these models, but we expect operators will find new ones. We look forward to working with operators to try specific analyses to support their operations improvements.

3. Enhancements

The models presented in this paper can be extended for other operations assumptions simply, and the methods provided will work for them as well. Some envisioned enhancements are as follows.

- Add the ability to fail while in the first proactive state.
- Decouple the proactive and reactive models so they are separate, can overlap, and can be solved generally; then combine them with the degradation model.
- Add the ability to model two types of repairs after center triage, so that two different repair response processes can be modeled accurately; this adjustment is depicted in Figure 26.

These models were built without working with specific operators' parameters, so we expect them to be a starting point only, with changes needed to better model specific operator situations. Certainly, the parameters may need to be updated, but there may be operations processes used by some operators that are significantly different and need further changes to properly model the situation.

Nonetheless, the results presented may apply and inform many existing operators already, making it easier to justify trying proactive methods in their operations with greater confidence.

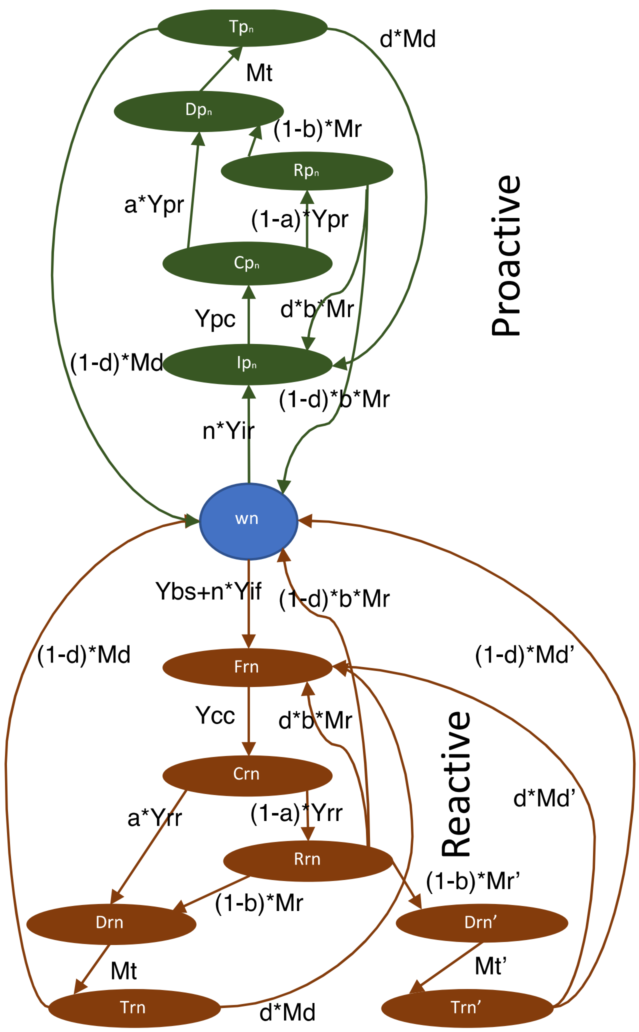

Figure 26 – An adjustment to the 11 state repair model to allow for two types of reactive repair handling, as recommended by an operator.

Abbreviations

CableLabs	Cable Television Laboratories
CM	cable modem
DOCSIS	Data-Over-Cable Service Interface Specifications
ISBE	International Society of Broadband Experts
PNM	proactive network maintenance
RF	radio frequency
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

[1] Taylor and Karlan “An Introduction to Stochastic Modeling,” Academic Press, 1984, ISBN-10: 0126848807.

[2] Jason Rupe, “Optimal Maintenance Modeling on Finite Time with Technology Replacement and Changing Repair Costs,” Annual Reliability and Maintainability Symposium. 2000 Proceedings. International Symposium on Product Quality and Integrity (Cat. No.00CH37055)