

Reliability-centered Maintenance of the Electrically Insulated Railway Joint via Fault Tree Analysis: A practical experience report

Enno Ruijters and Dennis Guck
University of Twente
Formal Methods and Tools
P.O. Box 217
7500 AE Enschede
The Netherlands
{e.j.j.ruijters, d.guck}@utwente.nl

Martijn van Noort
ProRail
P.O. Box 2038
3500 GA Utrecht
The Netherlands
martijn.vannoort@prorail.nl

Mariëlle Stoelinga
University of Twente
Formal Methods and Tools
P.O. Box 217
7500 AE Enschede
The Netherlands
m.i.a.stoelinga@utwente.nl

Abstract—Maintenance is an important way to increase system dependability: timely inspections, repairs and renewals can significantly increase a system’s reliability, availability and life time. At the same time, maintenance incurs costs and planned down time. Thus, good maintenance planning has to balance between these factors.

In this paper, we study the effect of different maintenance strategies on the electrically insulated railway joint (EI-joint), a critical asset in railroad tracks for train detection, and a relative frequent cause for train disruptions. Together with experts in maintenance engineering, we have modeled the EI-joint as a fault maintenance tree (FMT), i.e. a fault tree augmented with maintenance aspects. We show how complex maintenance concepts, such as condition-based maintenance with periodic inspections, are naturally modeled by FMTs, and how several key performance indicators, such as the system reliability, number of failures, and costs, can be analysed.

The faithfulness of quantitative analyses heavily depend on the accuracy of the parameter values in the models. Here, we have been in the unique situation that extensive data could be collected, both from incident registration databases, as well as from interviews with domain experts from several companies. This made that we could construct a model that faithfully predicts the expected number of failures at system level.

Our analysis shows that that the current maintenance policy is close to cost-optimal. It is possible to increase joint reliability, e.g. by performing more inspections, but the additional maintenance costs outweigh the reduced cost of failures.

I. INTRODUCTION

Reliability-centred maintenance (RCM) [1] is an important trend in infrastructural asset management. Its goal is to obtain optimal maintenance policies by maintaining crucial objects more intensively than less crucial ones. Thus, RCM tries to find an optimal balance between maintenance cost and system dependability, by placing maintenance effort where it matters most. To make such decisions, RCM requires a good insight in the effect of a maintenance policy on the system dependability, with key performance indicators as the system reliability, availability, and mean time between failures, etc. In fact, since RCM intertwines dependability and maintenance, it asks for an integral analysis of these two aspects. This paper demonstrates how such integral analysis can work and leads to useful results on RCM strategies, by studying a typical infrastructural asset via fault-maintenance trees.




Fig. 1. An electrically insulated joint with the visible components indicated.

Fault tree analysis (FTA) [2] is a popular methodology for dependability analysis. When the failure rates of the components are known, then FTA can compute the odds of a failure of the entire system. In practice, however, these failure rates are strongly affected by maintenance, which is not taken into account by fault trees. Thus, FTA is not suitable when the maintenance policy is subject to variation.

To overcome this limitation, and assess the impact of different maintenance strategies on system reliability and costs, fault maintenance trees (FMTs) have been developed [3]. These combine fault trees with *maintenance models*, representing the required ingredients for maintenance: component degradation, inspections, and repairs.

Moreover, FMTs necessitate the introduction of a new gate: the RDEP (rate dependency) gate makes that the failure of one component can accelerate the degeneration of other components. In this paper, we show that RDEPs are essential to faithfully model the EI-joint.

FMTs support the calculation of a number of important dependability metrics, such as the system reliability, availability, MTTF, expected cost etc. Technically, these analyses are realized via statistical model checking [4], a novel Monte Carlo

simulation technique [5].

EI-joints. Electrically insulated joints (EI-joints, see Figure 1) are an important railroad element, facilitating train detection and protection by electrically separating different track sections. They are a relatively frequent cause for failures and service disruption, so good maintenance is crucial for EI-joints.

Moreover, maintenance of the EI-joint is typical for other assets as well, with both random and wear-induced failures, repairs and renewals, and different options for maintenance strategies, and significant costs for failures and maintenance.

Modeling and analysis. In close collaboration with the Dutch national railway network infrastructure manager ProRail, we have conducted a reliability analysis of electrically insulated joints. We analyze the dependability of these joints, computing the reliability, expected number of failures, and expected costs over time. In particular, we investigate a reference maintenance strategy, as well as potentially better strategies. We study (1) variations in inspection intervals, (2) periodic preventive replacements, (3) replacement of an entire joint instead of repairing individual components, and (4) repairs when observing higher or lower degradation levels.

Our analysis finds that (1) the current inspection policy is nearly cost-optimal when combining cost of failure and cost of maintenance, (2) periodic preventive replacements improve reliability, but are more expensive than corrective replacements, and (3) the optimal inspection policy does not vary much with the load level of the track.

An important contribution is the extensive validation of our model: To provide confidence in the results of our analysis, we have compared the results predicted from our analysis with actual data from a failure database. Our predicted results agree with actual results from the field strongly enough to make recommendations based on our model.

Last but not least, we conclude that FMTs are a useful framework to investigate maintenance optimization problems from industrial practice: FMTs are a convenient model, have sufficient expressive power to capture complex maintenance aspects; and are able to produce predictive analysis results.

Related work. Many analysis techniques and extensions for fault trees exist, for an overview we refer the reader to [6]. Current FTA techniques support simple repair strategies by either equipping leaves with repair times [2] or with repair boxes [7], but do not consider preventive maintenance.

More complex repair policies are supported by the Repairable Fault Tree [8] formalism by Codetta-Raiteri et al., but this formalism still requires exponentially distributed failure and repair times.

Non-exponential failure time distributions can be used in the tool by Bucci et al. [9], which can be used to analyze component failures due to wear over time. This tool, however, does not consider maintenance to undo this wear.

Degraded states can be modeled in Extended Fault Trees by Buchacker et al. [11] which also supports components with failure rates that depend on the states of other components. Failure times are still modeled as exponential distributions, and this method does not include repairs or inspections dependent on full subtrees.

Looking outside FTA, Carnevali et al. [10] consider maintenance in phased systems where resources are used in a sequence of tasks, with detection and repair actions inbetween




Fig. 2. Depiction of the track circuit for train detection: the detection signal, depicted as the green line, is generated at the left of the images, and the detector is at the right. The top image depicts the situation where the track is clear, the red lines on the bottom picture indicate the axles of a train.¹

these tasks.

In systems consisting of identical components, Van Noortwijk and Frangopol [12] consider in detail two models of the effects of various maintenance choices on the reliability and cost in civil infrastructure, but these do not generalize to systems of multiple different components.

Organization of the paper. This paper begins with a description of EI-joints in Section II and the methodology in Section III. The modeling of the EI-joint by FMTs is explained in Section IV. Section V explains how this model is analyzed, and provides the results of this analysis. Finally, we provide our conclusions in Section VI.

II. CASE DESCRIPTION: MAINTENANCE OF EI-JOINTS

Electrically insulated joints (see Figures 1 and 2) are a railway component used in the detection of the occupancy of a railroad segment. They consist of a piece of insulating material between the ends of two tracks, to keep different segments of track electrically separated, while mechanically holding the tracks together.

Due to the large number of these joints in the railroad network, EI-joints are a relatively frequent cause of disruptions. Failures can occur for various reasons, both internal to the joint such as broken bolts, and external to the joint such as metal shavings bypassing the insulation. Inspections can be performed to determine whether some of these failures are likely to occur soon, and corrective action, such as sweeping away iron shavings, can prevent certain failures from occurring. Other failures can only be prevented or corrected by replacing the entire joint. Some failures, such as vandalism, cannot be prevented by maintenance.

A. Purpose and operation

Many railroad networks use electrical detection to determine the presence of trains on the tracks (e.g. in The Netherlands [13]). This system works by detecting when the axles of a train electrically connect the two rails, illustrated in Figure 2. To determine the location of a train, tracks are divided into several, electrically isolated, sections.

To detect the presence of a train, a small detection voltage is applied across the rails at one end of a section, and detected

¹Source: https://commons.wikimedia.org/wiki/File:Track_circuit.png

at the other end. A train on the section will short circuit the detection current, so the signal is not detected, and the interlocking system locks switches in their positions, sets signals appropriately, etc.

The location of a train is determined by creating electrically separate sections of track, each of which has its own detection current and detectors. On straight stretches of rail, these sections are several hundred metres to several kilometres in length. In areas with switches or level crossings, the sections are often much shorter.

B. Joint construction

The electrically insulated joint consists of a layer of insulating material placed between two sections of rail. The section of insulating material is called the *end post*. In glued joints, this post is produced at the factory attached to the ends of the rails, and the entire assembly including several metres of track is welded in place. In constructed joints, the end post is a separate component and mechanically held in place after assembly on site.

The rails are further held together by attaching one *fishplate* on each side of the rail with bolts. Insulating material is used to prevent the fishplates making electrical contact with the rails. Likewise, insulating bushings maintain separation between the bolts and the rails. Since the joint forms a weak point in the rail, two sleepers are normally placed close together where the joint is located, providing increased support to prevent the joint from flexing and breaking.

C. Failure modes

EI-joints are subject to two general categories of failures: *Mechanical failures* where the joint no longer provides a physical connection of the rails, and *electrical failures* that lead to an unintended electrical connection between the rails. The former type are uncommon, but have potentially catastrophic consequences (derailment of trains). The latter failures are more common and are generally not considered safety-critical due to the fail-safe nature of the detection system.

Table I lists the most significant failure modes, together with important failure parameters: Each mode is characterized by the expected time to failure assuming no maintenance is performed, the number of degradation phases we consider our modeling, and the probability that a given joint is subject to this failure mode. The latter is needed, since not all failure modes occur in all situations. For instance, Line 1 in the table shows that only 10% of the EI-joints are subject to poor geometry; 90% of the joints have a sufficiently stable surface to that this failure mode never occurs.

D. Inspections and repairs

A possible maintenance policy described by ProRail consists of several annual inspections, followed by corrective maintenance to repair any faults found by the inspection. This policy is taken as the reference policy in this paper.

The corrective action to be taken depends on the type of fault. Some faults, such as metal shavings causing a short circuit, can be immediately repaired without affecting any other failure mode. Other failure modes require a more general corrective action, such as grinding the surface of the rails, that also repairs wear of other failure modes. Finally, some failures require a complete replacement of the joint, thus repairing degradation of all other failure modes.

BE nr.	Failure mode	ETTF (yrs)	Phases	Prob. cnd.
1	Poor geometry	5	4	10%
2	Broken fishplate	8	4	33%
3	Broken bolts	15	4	33%
4	Rail head broken out	10	4	33%
5	Glue connection broken	10	4	33%
5a	Manufacturing defect	-	-	0.25%
5b	Installation error	-	-	0.25%
6	Battered head	20	4	5%
7	Arc damage	5	3	0.2%
8	End post broken out	7	3	33%
9	Joint bypassed: overhang	5	4	100%
10a	Joint shorted: shavings (normal)	1	4	12%
10b	Joint shorted: shavings (coated)	10	4	3%
11	Joint shorted: splinters	200	1	100%
12	Joint shorted: foreign object	250	1	100%
13	Joint shorted: shavings (grinding)	5000	1	100%
14	Sleeper shifted	5000	1	100%
15	Internal insulation failure	5000	1	100%
16	End post jutting out	20	1	100%

TABLE I. PARAMETERS OF THE BASIC EVENTS OF THE FMT FOR THE EI-JOINT. THE COLUMN ‘ETTF’ LISTS THE EXPECTED TIME TO FAILURE, ASSUMING NO MAINTENANCE IS PERFORMED. THE COLUMN ‘PROB. CND.’ GIVES THE PROBABILITY THAT A GIVEN JOINT IS SUBJECT TO THE CONDITION THAT ALLOWS THIS FAILURE MODE TO OCCUR. MODES 5A AND 5B HAVE A FIXED PROBABILITY OF OCCURRING EVERY TIME A JOINT IS INSTALLED.

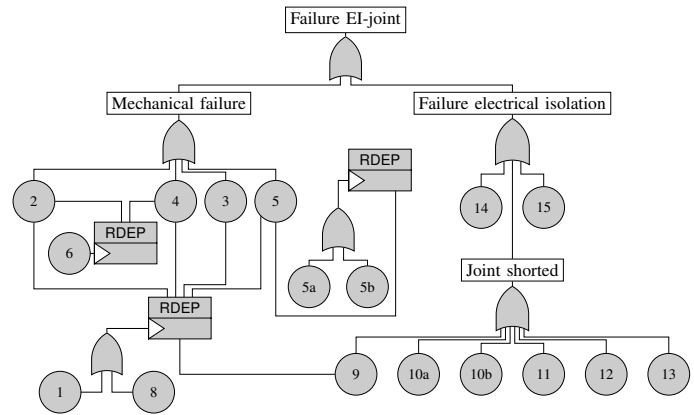


Fig. 3. Fault Tree describing the major failure modes of the EI-joint. The numbers in the basic events correspond to the section numbers of the failure modes. Failure modes 5a and 5b are specific causes of failure mode 5 (broken glue connection), due to manufacturing defects and installation errors, respectively. Failure modes 6, 7, and 16 have been merged into mode 6, as these are specific causes of the same fault. Failure mode 10 (short due to shavings) is separated into 10a for joints without additional protective coating, and 10b for joints with protective coating.

E. Problem Statement

We like to use the EI-joint to find out if FMTs are a useful tool to investigate maintenance questions, and to obtain trustworthy results. In particular, we like to know if the modeling power is sufficient to model the complex maintenance policies used in practice; if we can analyze relevant questions, and if we get faithful results that are useable in practice. The key question to be analysed for EI-joint is if the current maintenance strategy is effective and efficient. That is, whether the desired reliability requirements are met, whether it is cost effective, and whether improvements are possible.

III. METHODOLOGY

We have modeled the EI-joint in terms of fault maintenance trees. Below, we briefly describe the main ingredients of this framework: fault trees, maintenance models, analysis methods and metrics.

A. Fault Trees

Fault trees (FTs) are a widely used, graphical method for performing reliability and safety analysis [2] [6]. They are directed acyclic graphs in which the leaves, called *basic events* (BEs) describe component failures, and internal nodes, called *gates* or *intermediate events*, describe how these component failures interact and propagate to cause system failures. The root of the tree, called the *top level event*, denotes such system failure.

The gates of standard fault trees are AND-, OR-, and VOT(k)-gates, which fail when, respectively, all, any, or at least k of their children fail. The leaves are traditionally equipped either with failure probabilities, describing the probability of each leaf failing within the time of interest, or exponential failure rates, describing the progression of failure probabilities over time.

B. Fault maintenance trees

Fault maintenance trees (FMTs) [3] are an extension of FTs that can model several additional contributors to system reliability, including gradual degradation of components over time, inspections and repairs, and dependencies where one event triggers an accelerated degradation of another component. The FMT modeling the EI-joint is shown in Figure 3.

Extended basic events. The BEs in an FMT are more expressive than in standard BEs: Standard BEs usually model only failure or normal operation, with specific distributions of failure times such as exponential or Weibull distributions. An extended BE can be equipped with multiple phases, representing different stages of degradation. The transition into a next phase is described by an exponential distribution. Since the BE progresses linearly through the stages, the total failure behaviour of a BE in a FMT is described by an Erlang distribution.

RDEP gates. FMTs contain all the gates of static and dynamic FTs. Additionally, they contain a rate dependency (RDEP) gate, representing dependencies between components leading to accelerated wear. This gate has one trigger input, and one or more dependent children. When the trigger input fails, the failure behaviours of the dependent children are all accelerated by a factor γ , which can be different for each child. When the trigger input is repaired, degradation of the dependent children returns to their normal rate.

Repair and inspection modules. Standard FTs can support relatively simple repairs using distributions over repair times, or via repair boxes [7]. FMTs model more advanced maintenance policies via inspection and repair modules (IMs and RMs).

The IM describes at what frequency components are inspected as well as so called repair threshold. The latter is the (minimal) degradation phase where repairs will be performed. At degradation phases lower than the threshold, no repair will take place, either because the degradation is not visible, or because it is not considered necessary. When the threshold is passed, and the next inspection will trigger a repair. Thus, the IM will send out a repair request to the appropriate RM.

The RM listens for repair requests for the components under its control and initiates their repair or replacement. After the RM is invoked, the BE changes its phase to a less degraded

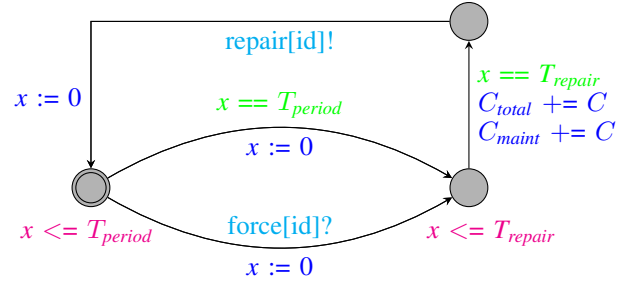


Fig. 4. PTA for a repair module. The PTA begins in the leftmost state with clock x initially zero. It waits until either the waiting time for a periodic repair (T_{period}) elapses, or a repair request signal ($force[id]$) is received. In either case, the module waits some time T_{repair} , incurs the C for a repair, sends a signal ($repair[id]$) so any BEs repaired by this module, and resets the timer.

phase. Moreover, the RM can invoke a periodic renewal of components, e.g. the replacement of a tire after four years.

C. Analysis of FMT by statistical model checking of priced timed automata

Technically, FMTs are realized via statistical model checking of price timed automata. That is, we first convert the FMTs into a network of priced timed automata (PTAs) [14] and use the statistical model checker Uppaal [15] to compute the relevant dependability metrics. Each element of the FMT (that is, each gate, BE, IM and RM) is translated into a price timed automaton. Then, all PTAs are composed together and analysed by Uppaal. We use the statistical engine here which is, unlike the verification engine, based on Monte Carlo simulation techniques.

PTAs are an extension of timed automata with costs on locations and actions. PTAs are transition systems that use real-valued clocks to specify deadlines and enabling conditions for actions. Costs can be incurred either with fixed amount when taking a transition, or by spending time in location, with a rate that is proportional to amount of time spent in the certain location.

The PTA for the repair module, inspection module and basic event are shown in Figures 4, 5, and 6, respectively.

During the translation, each FMT element (i.e. BE, gate, IM, and RM) is assigned a unique ID. The structure of the tree is then represented by the ids of the various signals used by the components to communicate. For example, if an IM with ID ‘1’ is inspecting a BE, the PTA for this BE will emit a signal $thres[1]$, to which the IM will react. The gates, not shown in this paper, listen for signals $fail[child_id]$ from their children, and emit their own signal $fail[id]$ when appropriate for their gate type.

D. Metrics

We analyze several aspects of the dependability of the EI-joint, which can be used to compare different maintenance policies and help in deciding which policy is better. We consider the reliability, expected number of failures, and costs.

Reliability. The probability of experiencing no system failures within a given time period. We compute the probability that within a certain period, there is never a time where a set of BEs is in a failed state leading to the occurrence of the top level event of the FMT.

Expected number of failures. We compute the expected number of occurrences of the top event in a given time window.




Fig. 5. PTA for an inspection module. The PTA begins in the leftmost state, and waits until either the time until the inspection interval (T_{period}) elapses, or until a threshold signal ($thres[id]$) is received from a BE. If the time elapses before a signal is received, then the inspection cost is incurred and the timer resets. If a threshold signal is received, the module waits for the scheduled inspection time, then signals its associated repair module to begin a repair ($force[rep_id]$), and then resets the timer.

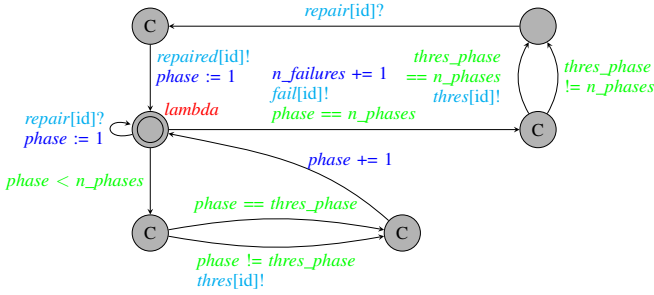


Fig. 6. PTA of a basic event with failure time given by an Erlang distribution with n_phases phases and an inspection threshold at $thres_phase$. From the initial state, the PTA waits an exponentially distributed time with mean $lambda$, and moves downward if it has not yet reached the last phase in the Erlang distribution, or rightward if it has. If it is not in the final phase, it advances by one phase, and it may emit a signal $thres[id]$ to a listening inspection module. The BE may also receive a signal $repair[id]$ and return to the initial phase. Upon completing the final phase, the failure counter is incremented and a signal $fail[id]$ is emitted. A threshold signal may be sent, and then the BE waits to receive a $repair[id]$ signal. After receiving this signal, the failed BE emits a signal $repaired[id]$, and returns to the initial phase and state.

Since all failures of the EI-joint can be repaired, there can be multiple failures over time. We can also compute the number of failures of individual components or subtrees of the FMT.

Cost. We can measure several costs incurred by the system over time. Specifically, we consider the costs of maintenance and failures. We can further separate costs into the costs of inspections, specific maintenance actions, and failures.

IV. MODELLING OF THE EI-JOINT

A. Fault tree modelling

The FMT has been constructed from a failure mode, effects, and criticality analysis (FMECA) [16] table that was provided by ProRail. An FMECA lists failure information per failure mode: its *effect* the consequences when this failure occurs, and its *criticality* describing how bad this failure is. In our case, the FMECAs are combined with, among others, the current maintenance policy, as well as failure frequencies.

The resulting FMT is displayed in Figure 3. As described in Section II-C, the joint failures are divided into physical and electrical failures. The electrical failures are further divided into failures caused by external influences such as iron shav-

ings short-circuiting the joint, and failures caused internally in the joint such as degradation of the insulating material.

The FMT for the EI-joint uses only ORs and RDEPs as gates. The method, however, works equally well with other FT gate types. The OR-gates show how to combine events into the top level event. The RDEPs are crucial to model failure dependencies, where the occurrence of one failure mode accelerates other failure modes. A few failure modes in the EI-joint have a severe effect on other failure modes: poor geometry affects almost all other physical failure modes; production and installation failures affect the failure of the glue connection, etc. Hence, a faithful model requires the expressive means to represent such failure accelerations.

The parameters of the BEs are listed in Table I.

B. Maintenance modelling

We compare the dependability and costs of joint subject to different maintenance policies. This allows us both to validate the model against actual recorded failure, and to offer suggestions for improvements in the policy that lead to cost savings or increased dependability.

ProRail has offered a possible maintenance policy, which is expected to reduce the number of failures to acceptable levels, and is close to the maintenance performed in practice.

In the FMT, inspection modules describe the inspection rates and the threshold at which corrective action is performed. The threshold in the FMT is described in terms of the degradation states of the BEs, while the reference policy describes physical observations such as ‘maximal vertical deformation 5 mm’. The translation of these physical descriptions to degradation phase was performed according to expert judgement.

Many BEs are maintained only by replacing the entire joint, which was implemented as a repair action that resets all BE degenerations to their initial state. The remaining BEs are maintained by correcting the specific fault identified during inspection, which is modelled by resetting only the degeneration of the BE undergoing the repair.

The current model makes a few assumptions: First, we assume that all inspections and repairs are carried out exactly on schedule. Since the fluctuations in inspection and repair times are small compared to the inspection interval, this assumption is reasonable. Also we assume that that inspections are perfect, i.e. an inspection always leads to a repair if the degradation level is past the threshold. While this may seem more questionable, we argue that the possibility of missing a failure is partially accounted for in the degradation threshold.

C. Choosing parameters for the model

One of the key factors in the analysis is the choice of the values for the parameters in our model. We have spent significant effort on the data collection process, via extensive consulting with domain experts at different contractors, leading to a model that provides enough confidence.

Our BE models contain the following parameters: (1) The number n of degradation phases, (2) The rate λ of the exponential distribution between these degradation phases. (3) The probability of the conditions that are necessary for failures to occur. (4) The maintenance thresholds, i.e. the minimum degradation level where maintenance is performed. (5) The acceleration rates for the RDEP gates.

We have estimated the values for these parameters by designing a questionnaire sent to several experts on mainte-




Fig. 7. Graph of degradation curves describing condition over time, provided as options for experts to describe degradation behaviour in the questionnaire.

nance for EI-joints. The appendix lists the questions asked. The responses from the maintenance experts mostly agreed. Further, we have used information from the aforementioned incident report system at ProRail.

Note that (1) and (2) together describe the time to failure for a given BE as an Erlang(λ, k) distribution. The expectation of this distribution equals λk , which should be equal to the expected life span L of the component if no maintenance is performed. The failure rates were directly asked in the questionnaire. The number of degradation phases were derived from the answer by the expert which failure curve shown in Figure 7 applies to each failure mode.

Certain failure modes can only occur if a certain condition exists. This condition is documented in the FMECA, and the probability of the condition per joint was obtained by an informal interview with an expert.

The acceleration rates are obtained from the FMECA, by comparing the indicated number of failures due to an RDEP-triggering failure to the total number of failures.

Having obtained these parameters, we have in some cases and together with experts further tuned the model, so that for each failure mode, the number of failures predicted by our model for that BE corresponds to the actual number of failures from the failure database.

Then, to validate our models, we have computed the number of maintenance actions required, and the total number of failures in one year. These values agreed with historical data recorded by ProRail, leading us to the conclusion that the tuned parameters are accurate.

D. Costs

Our model contains three categories of costs: *Failure costs*, *inspection costs*, and *repair costs*. To maintain confidentiality, the actual costs have been somewhat modified and no exact figures are shown in this paper.

Inspection costs are set as a fixed amount per inspection, and repair costs are fixed for each type of repair. The cost for failures consist of the cost caused by the unavailability of the railroad tracks. These are defined as societal cost, i.e., a synthetic costs that are used as a key performance indicator to steer the performance of railroad companies. These societal costs are also incurred when the tracks are unavailable due to planned maintenance.

V. ANALYSIS AND RESULTS

In this section we describe the results of several experiments we conducted on the FMT of the EI-Joint. As a first step, we validated the FMT against observations from the field. Therefore, we used the model as constructed, i.e. we analysed the EI-joint under the current policy. Since we concluded that the model is in line with the real world, we continued with

BE	Failure cause	Predicted	Actual	Difference
1	Poor geometry	110	48	62
2	Broken fishplate	129	83	46
3	Broken bolts	2.3	2.1	0.2
4	Rail head broken out	68	30	38
5	Glue connection broken	70	37	33
6	Battered head	3.4	5.5	2.1
7	Arc damage	7	3.4	3.6
8	End post broken out	12	9.4	2.6
9	Joint bypassed: overhang	212	200	12
10	Joint shorted: shavings	156	150	6
11	Joint shorted: splinters	254	261	7
12	Joint shorted: foreign object	199	200	1
13	Joint shorted: shaving from grinding	10	10	0
14	Damage by maintenance	19	18	1

TABLE II. COMPARISON OF PREDICTED AND ACTUAL FAILURE RATES OF DIFFERENT FAILURE MODES. VALUES ARE YEARLY OCCURRENCES IN A POPULATION OF 50,000 JOINTS.

finding possible improvements of the current policy. Therefore, the maintenance strategy within the FMT was modified by changing inspection frequencies and replacements. This led to a description of how an optimal maintenance strategy of the EI-Joint can be constructed.

Note that the results in this section are averages of 40,000 simulation runs each. The variance between the simulation runs is low enough that a 95% confidence interval around the mean results has a width less than 1% of the indicated value.

A. Current policy

First, we estimate the total failure rate of the joint over time, shown in Figure 10. This number is within the margin of error of ProRail's incident tracking. We further note that after approximately two years, the expected number of failures per year is almost constant.

To validate the model, the expected number of occurrences of each failure mode per year was estimated. Table II shows the predicted and actual number of occurrences of each failure mode. Note that ProRail maintains a record of joint failures by cause, and we compare the predicted number of failures to the recorded number. Since the predicted failure rate is almost constant, we assume we can multiply the expected failure rate by the number of joints to obtain the total number of failures, regardless of the age of the joints in operation. A graphical breakdown of the causes of failures is displayed in Figure 8.

The difference between actual and predicted failure rates for BE 1 is likely explained by inaccurate reporting, as engineers often report only the immediate defect rather than the underlying poor geometry. BEs 2, 4, and 5 concern mechanical failures, which are typically often corrected during maintenance before the officially specified threshold is reached.

As an additional validation, we estimate how often a joint is replaced due to maintenance. Our model predicts approx. 3680 replacements per year, on a population of 50,000 joints.




Fig. 8. Breakdown of failures of the EI-joint by cause. The numbers in the bottom row indicate individual failure modes, and correspond to the numbers in Table I.

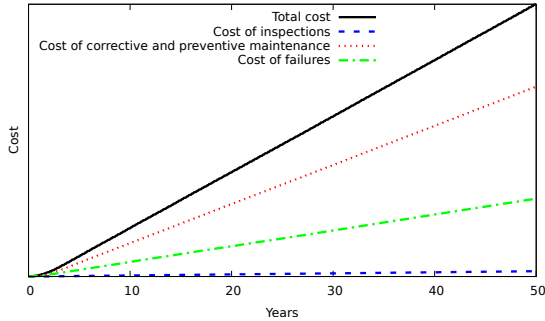


Fig. 9. Cumulative costs of one EI-joint over time, split up by type of cost.

ProRail records indicate approx. 3000 replacement joints are installed each year. We expect that this difference is due to some failure modes where the maintenance action induces a replacement in the model, whereas in some cases in the real system the degradation may not have progressed so far, resulting in only a minor maintenance action.

Next, we consider the costs of the joint. Figure 9 shows the various costs over the lifetime of the joint. As can be expected from the progression of the cumulative number of failures, also the costs progress very linearly over time. Although these numbers are fictionalized, the actual values do not deviate much from ProRail's estimate.

B. Optimization of maintenance policy

Having concluded that the model is a reasonably accurate description of the behaviour of the EI-joint, we present some options for improving the reliability and/or costs of the joint.

Inspection frequencies. First, we consider the possibility of performing more or fewer inspections. Figure 10 shows the cumulative expected number of failures over time for different numbers of inspections. We note that the introduction of any inspections at all significantly reduces the number of failures, but subsequent increases of the number of inspections have a much smaller effect. This is due to failures either occurring gradually and being detected even with infrequent inspections, or occurring suddenly, and rarely being found by any inspection before failing.

In terms of improving reliability, clearly more inspections are always better. Nonetheless, these results show diminishing returns when increasing the inspection frequency above approximately two per year.

To estimate the cost-optimal number of inspections, we plot the total cost per year for different inspection frequencies,

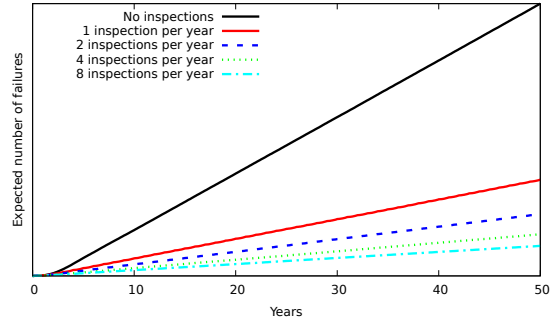


Fig. 10. Cumulative expected number of failures of one EI-joint over time, for different inspection rates.

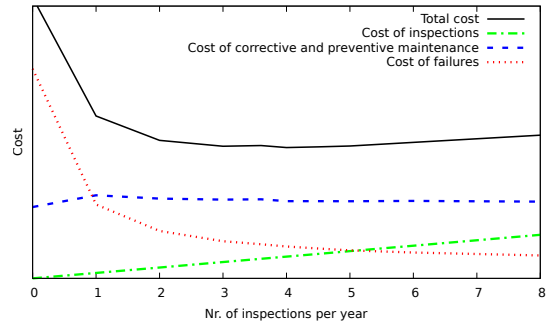


Fig. 11. Different types of total costs for one joint, depending on the inspection frequency.

shown in Figure 11. As expected, the costs of failures decrease with more inspections, while the costs of inspections increase. The maintenance costs are fairly constant, as increased inspections do lead to more necessary repairs, only repairs performed sooner.

The optimal number of inspections in terms of total cost is found around four inspections per year. The difference in total cost between approx. 2 and 6 inspections per year falls within the margin or error of the simulation, so no more precise optimum can be determined.

Replacements. Several other options for maintenance policies are listed in Table III. We consider always replacing the entire joint when any maintenance is required, adjusting the inspections to take preventive action well before the reference threshold, and periodically replacing the joint regardless of inspection result. We again find that all these policies have higher total cost than the reference policy. The reduced threshold on inspections does significantly decrease failures for only a modest increase in total cost, but since total cost includes the social cost of failure, we do not consider this a net gain. It is also questionable whether all failure modes show signs of wear sufficiently early to allow this policy to be implemented.

It is likely that the failure rates of the joint vary depending on the intensity of their use. Additionally, costs of unavailability due to failure or repair increase as the number of passengers passing over the joint increases. We have not precisely determined the correlation of these effects, but we have analysed the optimal inspection frequency for several variations of costs and failure rates. The optimal inspection frequencies are listed in Table IV, as well as the relative cost of the optimal inspection

Policy	Maint. cost	Total cost	Failure frequency
Current	1	1	1
Replace instead of repair	2.20	1.65	0.76
Reduce threshold by $\frac{1}{3}$	1.49	1.16	0.48
Replace every 5 yrs.	2.49	1.85	0.88
Replace every 10 yrs.	1.59	1.34	0.96
Replace every 20 yrs.	1.30	1.17	0.97

TABLE III. COMPARISON OF THE EFFECTS OF DIFFERENT MAINTENANCE POLICIES, RELATIVE TO THE REFERENCE POLICY.

Optimum Cost factor	Failure rate factor			Rel. cost Cost factor	Failure rate factor				
	2	$\frac{3}{2}$	1		$\frac{1}{2}$	2	$\frac{3}{2}$	1	$\frac{1}{2}$
$\frac{1}{2}$	8	8	5	2	$\frac{1}{2}$	0.94	0.99	0.98	0.91
1	8	8	4	2	1	0.92	0.99	1	0.92
$\frac{3}{2}$	8	6	4	2	$\frac{3}{2}$	0.92	0.96	1	0.89
2	6	6	3	2	2	0.94	0.98	0.98	0.88

TABLE IV. OPTIMAL INSPECTION FREQUENCIES PER YEAR FOR DIFFERENT RELATIVE FAILURE RATES AND COSTS, AND TOTAL COST OF THIS POLICY COMPARED TO THE REFERENCE POLICY (4 PER YEAR). ALL COSTS (I.E. INSPECTION, REPAIR, AND FAILURE) ARE INCREASED BY THE SAME FACTOR.

policy compared to the previously computed optimum of 4 inspections per year.

We find that the optimal inspection frequency is determined primarily by the degeneration rate, rather than by the cost. Furthermore, the optimal inspection policy has at most a 12 percent cost saving compared to a general policy of four inspections per year.

VI. CONCLUSION

We have modeled and analyzed several maintenance policies for the EI-joint via fault maintenance trees. We conclude that obtaining the FMT for the EI-joint was not too difficult from the information in the existing FMECA. Obtaining the right quantitative information required additional effort, but was feasible as well. We found that FMTs naturally model the EI-joint, and is a useful tool to investigate different maintenance policies.

One may wonder how surprising it is that the reference maintenance strategy is cost optimal under the existing circumstances. We argue that it might not be so, because the EI-joint is a well-understood railroad element. Nevertheless, our analysis has provided useful insights in the degradation behavior of the joints, for instance in critical accelerating factors.

Future work includes the extension of FMTs with continuous degradation phases, models that take into account specific conditions and usage scenarios that influence degradation. Additional work could include different analysis techniques such as rare-event simulation or analytic approaches that could allow FMTs to be used for systems where highly improbable events have significant effects.

ACKNOWLEDGMENT

This work has been supported by the STW-ProRail partnership program ExploRail under the project ArRangeer (122238) with participation by Movares. We thank Judi Romijn and Jelte Bos for their helpful comments on earlier versions of this paper.

REFERENCES

[1] J. Moubray, *Reliability centered maintenance*. Industrial Press, 1997.
[2] W. E. Vesely, F. F. Goldberg, N. H. Roberts, and D. F. Haasl, *Fault Tree Handbook*. U.S. Nuclear Regulatory Commission, 1981.

[3] E. Ruijters, D. Guck, P. Drolenga, and M. Stoelinga, "Fault maintenance trees: reliability centered maintenance via statistical model checking," in *Proc. of the Reliability and Maintainability Symposium (RAMS)*, 2016.
[4] A. Legay, B. Delahare, and S. Bensalem, "Statistical model checking: An overview," in *Proc. 1st Int. Conf. on Runtime Verification (RV)*, ser. LNCS, vol. 6418, Nov. 2010, pp. 122–135.
[5] G. Fishman, *Monte Carlo: Concepts, Algorithms, and Applications*. Springer, 1996.
[6] E. Ruijters and M. Stoelinga, "Fault tree analysis: A survey of the state-of-the-art in modeling, analysis and tools," *Computer Science Review*, vol. 15–16, pp. 29–62, 2015.
[7] A. Bobbio and D. Codetta-Raiteri, "Parametric fault trees with dynamic gates and repair boxes," in *Proc. 2004 Annual Reliability and Maintainability Symposium (RAMS)*, 2004, pp. 459–465.
[8] D. Codetta-Raiteri, G. Franceschinis, M. Iacono, and V. Vittorini, "Repairable fault tree for the automatic evaluation of repair policies," in *Int. Conf. Dependable Systems and Networks*, 2004, pp. 659–668.
[9] G. Bucci, L. Carnevali, and E. Vicario, "A tool supporting evaluation of non-markovian fault trees," in *Proc. 5th int. conf. on Quantitative Evaluation of Systems (QEST)*, Sep. 2008, pp. 115–116.
[10] L. Carnevali, M. Paolieri, K. Tadano, and E. Vicario, "Towards the quantitative evaluation of phased maintenance procedures using non-markovian regenerative analysis," in *Proc. 10th European Performance Engineering Workshop*, ser. LNCS, vol. 8168, Sep. 2013, pp. 176–190.
[11] K. Buchacker, "Modeling with extended fault trees," in *Proc. 5th IEEE International Symposium on High Assurance Systems Engineering (HASE)*, 2000, pp. 238–246.
[12] J. M. van Noortwijk and D. M. Frangopol, "Two probabilistic life-cycle maintenance models for deteriorating civil infrastructures," *Probabilistic Engineering Mechanics*, vol. 19, no. 4, pp. 345–359, Oct. 2004.
[13] ProRail, "Netverklaring 2016, gemengde net [in dutch]," 2015. [Online]. Available: <https://www.prorail.nl/vervoerders/netverklaring>
[14] G. Behrmann, K. G. Larsen, and J. I. Rasmussen, "Priced timed automata: Algorithms and applications," in *Formal Methods for Components and Objects*, ser. LNCS, vol. 3657, 2005, pp. 162 – 182.
[15] P. Bulychev, A. David, K. G. Larsen, M. Mikučionis, D. B. Poulsen, A. Legay, and Z. Wang, "UPPAAL-SMC: Statistical model checking for priced timed automata," in *Proc. 10th workshop on Quantitative Aspects of Programming Languages (QAPL 2012)*, 2012.
[16] M. Rausand and A. Hoylan, *System Reliability Theory. Models, Statistical methods, and Applications*. Wiley, 2004.

APPENDIX

To obtain information about the failure behaviour of the components of the EI-joints, a questionnaire was sent to several experts. The exact questions were:

- 1) What is the average time until this failure mode occurs, assuming no maintenance is performed?
- 2) Are there conditions that occur regularly and significantly affect the time to failure? If so, what are these conditions and what effect to they have on the time to failure?
- 3) Which of the graphs best describes the degeneration behaviour of this failure mode? [The graphs in Figure 7 were included.]
- 4) If an inspection is performed around half the expected time to failure, is it likely that clear signs of wear will be found?
- 5) If an inspection near the expected time to failure does not find indications of wear, is it likely this failure mode will occur much later than estimated?
- 6) Does this failure mode frequently occur shortly after installation?
- 7) How often does this failure mode occur before less than half the expected time has passed?
- 8) How often does this failure mode only occur later than 1.2 times the expected time?
- 9) How often does an inspection lead to a maintenance action?
- 10) If an inspection shows a need for maintenance, how soon after the inspection must this maintenance be performed to prevent failure?