The published version of this paper can be found at http://dx.doi.org/10.1109/RAMS.2016.7447986.

Fault maintenance trees: reliability centered maintenance via statistical model checking

Enno Ruijters, University of Twente

Dennis Guck, University of Twente

Peter Drolenga, NedTrain

Mariëlle Stoelinga, University of Twente

Key Words: Fault tree analysis, maintenance, repairable systems

SUMMARY & CONCLUSIONS

The current trend in infrastructural asset management is towards risk-based (a.k.a. reliability centered) maintenance, promising better performance at lower cost. By maintaining crucial components more intensively than less important ones, dependability increases while costs decrease.

This requires good insight into the effect of maintenance on the dependability and associated costs. To gain these insights, we propose a novel framework that integrates fault tree analysis with maintenance. We support a wide range of maintenance procedures and dependability measures, including the system reliability, availability, mean time to failure, as well as the maintenance and failure costs over time, split into different cost components.

Technically, our framework is realized via statistical model checking, a state-of-the-art tool for flexible modelling and simulation. Our compositional approach is flexible and extendible. We deploy our framework to two cases from industrial practice: insulated joints, and train compressors.

1 INTRODUCTION

Safety-critical systems, like medical devices, trains, airplanes and nuclear-power plants must be dependable. Stakeholders require substantiated facts and figures regarding system performance, costs and risks to support decision making and demonstrate compliance.

The aspect of safety and reliability is usually tackled by applying a Fault Tree Analysis (FTA) [1]. In contrast to other popular methods like failure mode and effect analysis (FMEA), which is a bottom-up technique to identify and prioritize potential failures of a system, FTA is a top-down failure analysis and focuses on discovering root causes of system failures.

Fault trees (FTs) describe how component failures propagate through a system. They consist of *gates*, modelling the failure propagation, and *leaves*, modelling component failures which are equipped with a probability distribution describing the component's failure behavior over time. It should be noted that these probabilities are conditional with respect to the execution of a maintenance program. That is, given the design, system usage, environmental context, maintenance program and human factors, residual probabilities are estimated for each of the root causes that ultimately may lead to one or more undesired top events.

FTA is mostly concerned with deriving critical failure behaviors in systems by determining the root causes. Although such an analysis can provide insight into the most likely failures, in itself it is not a tool to optimize for example repair policies, rejection levels and inspection intervals. To do so, the fault tree must be connected to a system model that comprises the dynamics of the physical condition and the impact of maintenance execution. This allows one to investigate the effect of different maintenance policies on the reliability and availability of a system in the framework of FTA.

This paper presents fault maintenance trees (FMTs), a formalism connecting traditional FTA with maintenance strategies. The key idea is to encode the degradation of components into the FT as well as different inspection regimes. We show how to implement preventive as well as corrective maintenance actions into FTA based on these extensions.

2 FAULT MAINTENANCE TREES

Current FTA techniques support repairs by equipping leaves with repair times [1] or repair boxes [2]. Maintenance, however, is much more complex, involving inspections, renewals, degradations, and many more. Such phenomena are not yet supported by fault trees. We propose an extended model to capture the physical system condition and a variety of compensating maintenance actions. An integrated model yields the effect on dependability measures as safety, reliability and availability.

We deploy dynamic fault trees (DFTs) as basis of our fault maintenance trees (FMTs). A DFT is a tree — or rather a directed acyclic graph — describing the failure propagation

throughout a system in terms of its components. DFT leaves model the components failure behavior (called *basic events*) and all other nodes (called *gates*) describe the failure propagation, where the root node is called the top level event (TLE).

The static OR-, AND-, and VOT(k)-gate fail if respectively, one, all or k of their inputs fail. The dynamic gates PAND, SPARE and FDEP already encode some common reliability patterns like sequencing, spare management and dependencies and are described in [3].

Basic events (BEs) represent the components' failure behavior over time. We take a standard approach and consider stochastic failure behavior modelled by exponential distributions. Thus, the probability that a component fails within time t is given by $P[X < t] = 1 - e^{-\lambda t}$. Apart from exponential distributions, acyclic phase-type distributions can be used to approximate any probability distribution with arbitrary precision.

Our fault maintenance trees (FMTs) augment DFTs in three ways. We introduce (1) maintenance models describing the effects of maintenance actions; (2) additional gates to model phenomena that occur w.r.t. maintenance; (3) different cost values for maintenance as well as for down time and repairs. Further, we analyze these FMTs w.r.t. various dependability metrics and costs.

2.1 Maintenance models

Maintenance models are specified in conjunction with the DFT and specify the effects of maintenance actions on the leaves. Therefore, we redefine the behavior of BEs and introduce inspection and repair modules.

Degradation. The condition of a component is modelled by different degradation phases in the leaves, similar to extended FTs [4]. For example, a new tire is considered as fully functional in the beginning, but by ageing as well as usage the quality degenerates, e.g. after 10000 miles the profile depletion is already visible. By including such phenomena as separate phases in the leaves, we are able to observe it during an inspection.

Maintainable BEs. Each maintainable BE (MBE) is divided in n phases, where n is the number of different condition stages of the component. Additionally, a threshold specifies at which phase an inspection should trigger a maintenance action. The transition into a next phase is described as in BEs by an exponential distribution. Thus the failure behavior of an MBE is described by an acyclic phase-type distribution. This allows us to approximate any arbitrary distribution [5].

Example 1. Figure 1 depicts an example distribution for an MBE. The bars represent the probability of a component failure for each year over a time span of 10 years, based on historical failure occurrences of the component. To approximate this failure behavior, we choose an Erlang(6,1.2) distribution. Hence, we have 6 degradation phases for the example component where each transition into the next phase is described by an exponential distribution with rate 1.2.

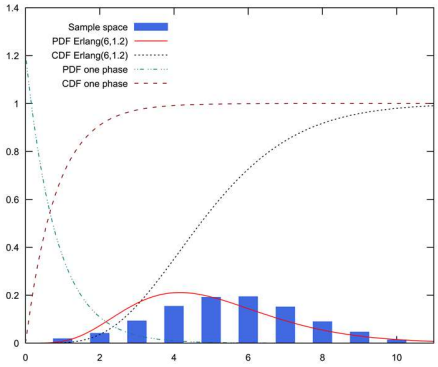

Figure 1 – Example of a MBE failure distribution.

Repairs and inspections. We distinguish between two types of maintenance procedures: (1) corrective, and (2) preventive maintenance.

Corrective maintenance is carried out after a component has failed, replacing or repairing the broken component. The detection of a failed component can be either instantaneous due to an immediate impact on the system performance, or it can be hidden, i.e. the system performance is not affected but the structure of the system is weakened. In the first case a repair can be scheduled directly, whereas in the second case an inspection will be required to detect the failure and initiate the repair. Besides, a failure of a component can also cause an overhaul of the whole system. Thus, not only the failed component will be repaired, but all components will be inspected and refurbished.

Preventive maintenance refers to actions that try to prevent a component's failure. Hence, components are inspected, and based on the inspection a repair or (partial) renewal will be performed, putting the component into a better condition. Preventive maintenance can be condition-, time- or usage-based, e.g. the replacement of a car tire when the profile is too low, or too old, or after 100.000 miles, respectively. To encode corrective and preventive maintenance procedures, we use repair models (RM) and inspection models (IM).

The RM listens for repair requests of components and initiates their repair or partial replacement. Thus, after the RM is invoked, the MBE changes its phase from failed or degraded to a less degraded phase. Further, the RM can invoke a periodic renewal of components, e.g. the replacement of a tire after four years.

The IM describes at what frequency components are inspected as well as at what degradation phase a repair request will be send out. Note that we support different probability distributions: degradation is usually modelled stochastically, e.g. by an exponential or Weibull distribution, whereas inspections are typically done at a fixed frequency.

Additional behavior. By including maintenance and the repair of failed components into the framework of DFTs, we are shifting the perspective from a simple bottom up failure propagation to a reversible failure propagation. Consider an AND-gate with two components. When the AND-gate receives a failure signal from both components, it will emit a

failure. However, if one component gets repaired, the failure of the AND-gate is reversed, which needs to be propagated by a new signal from the component to the gate and so forth.

Additional, there are other phenomena that occur w.r.t. maintenance. Concretely, we introduce a rate-dependency (RDEP) gate where a trigger (e.g., installation failure or bad coating) causes dependent events to fail at a higher rate. This is not only restricted to a maintenance actions and can also be used to describe dependable failure accelerations of components, e.g. if there is a leak in the cooling system of a motor, the failure rate of the motor will get accelerated.

2.2 Costs

The integration of costs is a crucial factor in the decision process of a maintenance strategy. We distinguish between two cost factors: (1) maintenance costs, including inspections and maintenance related repairs as well as overhauls and (2) failure costs, including the downtime costs and the replacement costs of broken components.

Those different cost factors can be incurred with a certain rate over time, or instantaneously with a fixed amount per action. For example, if the system is not operational, the costs will increase by each time unit until it is restored, whereas a replacement of a broken component will have a fixed cost. Additionally, the distinction between maintenance and failure costs allows for a cost-benefit analysis of different maintenance strategies.

3 METHODOLOGY

As depicted in Figure 2 our analysis is realized via compositional transformation [6] and statistical model checking: We model each FT element (i.e. gate, leaf, or maintenance specification) as a timed automaton (TA), reflecting the state-based behavior of the element: the TA models for the leaves contain the various degradation phases and failure rates. Inspections are modelled by a TA parametrized with the inspection frequency; similarly for repairs, we model the repair strategies (e.g. first come first serve) and repair times by a TA model. Each FT gate has an associated TA model that reflects how failures propagate. In this way, we obtain the TA model for the entire fault tree by composing the separate TA models. This yields one large TA which can be analyzed via statistical model checking. That is, we let the statistical model checker UPPAAL [7] generate a number of traces of the TA model, from which we estimate the metrics under consideration.

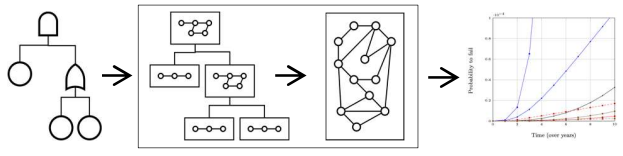

(a) FMT (b)Transformation (c) Model (d) Result

Figure 2 – The system is described as an FMT (a) where each element is translated into a TA (b). The separate TA models are analyzed as one whole model (c) and UPPAAL produces our desired. metrics (d).

3.1 Priced Timed Automata

We use priced timed automata (PTAs) [8], an extension of TAs with costs on locations and actions. A TA consists out of locations and edges. At any point in time, one location of the TA is the current location, and edges may be taken to make another location current. Constraints on the edges can be used to restrict the times when an edge may be taken, and invariants on locations prevent the TA from entering or remaining in that location at certain times. Constraints and invariants are specified in terms of clocks, which increase in value at a constant rate but may be reset when taking an edge. PTAs extend TAs with the addition of costs: Each location has an associated nonnegative cost which is incurred per unit time that the location is active. Similarly, edges have an associated cost that is incurred every time the edge is taken.

In standard PTAs, at any point in time an edge is either guaranteed to be taken (when the invariant on the current location is about to be violated), guaranteed not to be taken (when its constraint prevents it), or there is a nondeterministic choice whether to take it or not. The PTAs used in statistical model checking (SMC) have an additional option: locations can have an associated exponential rate, in which case an edge will be taken at a time determined by an exponential distribution. Furthermore, labels on edges can be used to require that edges in different PTAs be taken at the same time. Consider the PTA in Figure 3: The left location is the initial location, and *timer* is a clock. While this clock is below the constant *interval*, the invariant is true and the PTA may remain in this location. If the inspection threshold edge is taken in another PTA, the transition is taken to the rightmost location. Assuming this does not occur, when *timer* reaches the value *interval*, the invariant becomes violated and an outgoing edge must be taken. At the same time, the self-loop edge becomes permitted and thus is taken. Besides, this edge incurs some costs and resets the timer, thus it starts a new inspection period.

Example 2. Figure 3 and Figure 4 depict PTAs of an IM and an MBE. As described in Section 2.1 the MBE has different phases of degradation which change with rate λ_i. Note that the lightly degraded phase is split up into undetectable and detectable. This is important in the

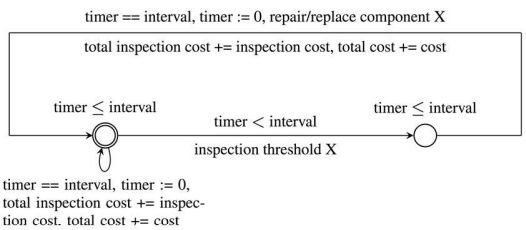

Figure 3 – PTA of an inspection module.

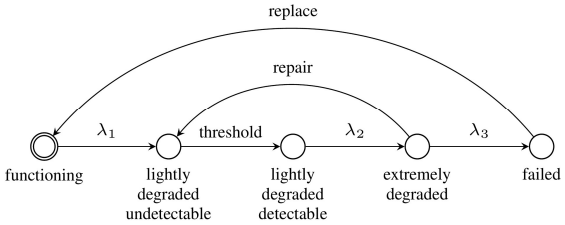

Figure 4 – PTA of a maintainable BE.

communication with the IM. The PTA of the IM has only two states, the first state describing an inspection with no findings, and the second state with findings, respectively. After the threshold signal from an MBE, the IM will induce a repair or replacement of when the interval is reached.

3.2 Statistical model checking

Statistical model checking provides multiple advantages over other simulation methods: it has rigorous mathematical foundations; we can validate our models by traditional, non-stochastic model checking and our compositional approach is easy to modify and extend.

We use the statistical model checker UPPAAL. The tool facilitates both a graphical and a textual user interface. The modelling of the basic FMT components is done with the graphical tool. The specification of the maintenance strategies, as well as the intercommunication in the FMT is described in textual format. The simulation allows us to examine the system by exploring a particular execution trace in each run. In contrast to the simulation, the model-checking part allows for a complete exploration of the system to check for invariant and reachability properties.

3.3 Quantitative Analysis

FMTs support a wide variety of common dependability metrics, including reliability, availability, and expected number of failures. In addition FMTs allow the analysis of expected costs, with the possibility to separate them into per-component costs as well as costs for downtime, repairs, inspections and more.

Dependability analysis. For our dependability analysis, we focus on two factors: (1) the reliability of the system, and (2) the expected number of failures.

The reliability is defined as the probability that the FMT has not failed until a given mission time $t \in \mathbb{R}_{>0}$, i.e. the TLE has not emitted a failure up to time point t.

The expected number of failures is given by the number of TLE failures during a give mission time t, i.e. the accumulated number of failures during a systems runtime. Those can also be broken down to component failures.

Cost analysis. The inclusion of costs into FMTs allows for several cost analyses. We focus on the computation of expected costs. This enables us to compute the probability of exceeding a certain budget. Further, by varying the parameters in the model, the influence of, e.g. the inspection frequencies, renewal and repair policies can be visualized w.r.t. the dependability.

4 CASE STUDY

We apply our framework to two industrial case studies from railway engineering. The first case study concerns an air compressor and was provided by NedTrain, subsidiary of the Nederlandse Spoorwegen (NS), responsible for rolling stock maintenance in the Netherlands. There are a total of 230 compressors of this type deployed in trains.

The second case study is the electrically insulated joint (EI-joint) and was provided by ProRail, the Dutch railroad asset management organization. EI-joints are part of a system for train detection and a frequent cause for delays and disruption in service. There are circa 45000 EI-joints in the Netherlands.

The actual data used in the case studies are anonymized.

4.1 Train compressor

The occurrence of stranded trains (downtime > 10 minutes during service) is an important KPI (key performance indicator) for NedTrain due to customer demands. The compressor itself provides the trains pneumatic system, used for brake and door control, with pressurized air. This is a critical system, as failures can lead to stranded trains.

Maintenance modelling. Figure 5 depicts the fault tree of one train compressor system. The events leading to a stranded train are divided into two event categories: (a) automatic safety actions, and (b) degradation of internal components.

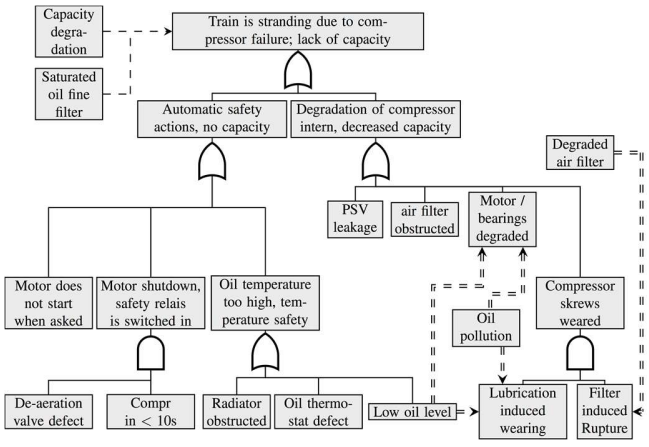

Figure 5 – Part of the FMT for the train compressor system. The simple dashed lines denote BEs which are used to model maintenance triggers, but cannot lead to a TLE failure. The double dashed lines denote dependencies which are used to model the change of failure rates on the dependent events (RDEPs).

The maintenance procedures for the train compressor can be divided into four categories: (1) daily inspections and minor corrective repairs at service locations near the track; (2) monthly check-ups with preventive and corrective maintenance at a depot; (3) small overhaul at the refurbishment and overhaul workshop every 3 years; (4) a complete overhaul of the compressor at the refurbishment and overhaul workshop after 6 years.

Results. Figure 7 shows the number of compressor failures under various maintenance strategies. It is clear that the 2-daily inspections, although they only find relatively obvious faults, strongly reduce the number of failures. Other variations on the maintenance policy do not have effects nearly as significant. Additional services are performed when the 2-daily inspections uncovers a defect. The variations on the other maintenance actions do not substantially affect the additional services. Comparing the results with the no-2-daily-inspection line of Figure 7 shows that the number of unscheduled services corresponds quite well to the number of failures prevented by this inspection.

4.2 Electrical insulated joint

An electrically insulated joint is a component that physically joins two sections of rail track, without creating an electrical connection between them. This allows the electrical train detection to discern the location of a passing train. Failures of these joints typically result in erroneous indications that trains are present in unoccupied sections of track, resulting in delays.

Maintenance modeling. Figure 6 shows the fault tree for one EI-joint. The failures are divided into two categories: failures that compromise the physical support of the rail, and failures of the electrical insulation. A failure in this case is considered any situation in which rail traffic is disrupted due to the joint, which usually occurs well before the situation actually becomes dangerous.

The maintenance policy is relatively straightforward: Inspections are carried out periodically, and repairs or replacements are performed when these inspections find problems. Since the type of joint studied here is provided as an integrated component by the manufacturer, many failures are repaired by replacing the entire joint. Some failures only affect the surface of the rail, and are corrected by grinding a thin layer off of the joint. The remaining failures are corrected by specific actions, e.g. removing foreign conductive material.

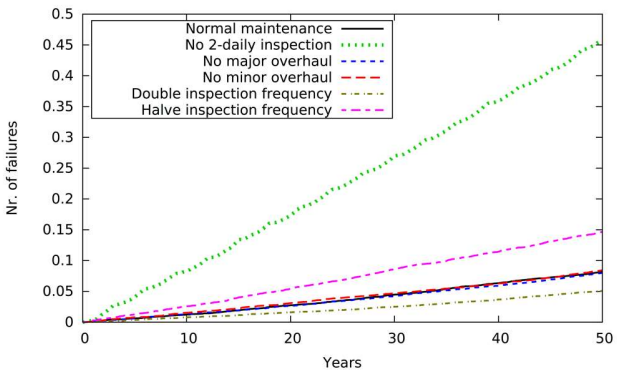

Figure 7 – Failures of a compressor under different maintenance policies.

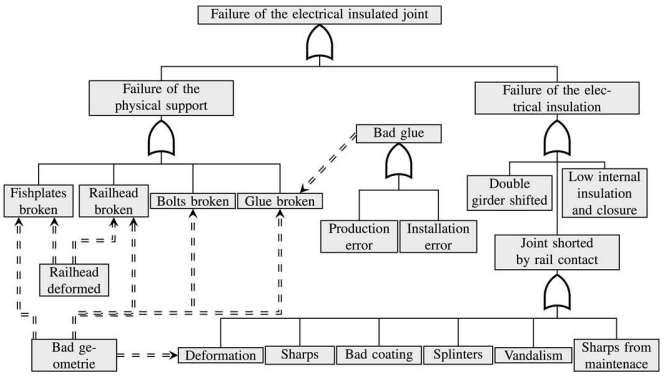

Figure 6 – Part of the FMT for the EI-Joint.

Results. Figure 8, Figure 9, and Figure 10 show the maintenance costs over time for the EI-joint.

Figure 8 depicts the cost breakdown for the current maintenance regime and shows that the total costs are dominated by the failure costs. Inspection and repair costs are very low. This suggests that investing in preventive-maintenance is cost-efficient. Figure 9 shows the cumulative costs for different inspection intervals. Finally, Figure 10 provides the total costs per inspection interval. We see that a cost-optimal inspection frequency is obtained with inspecting between 10 and 15 times per year, i.e. monthly inspections.

Note that the cost of an inspection is relatively difficult to specify. Most inspections are performed together with maintenance on other components, so the marginal cost of this inspection is very low. If additional inspections are performed only on the EI-joint, these may have substantially higher cost. If an inspection requires rail traffic to be stopped, the cost would be almost as high as the cost of a failure.

5 CONCLUSION AND FUTURE WORK

This paper presents a novel extension to FTA by applying maintenance directly on the fault tree level. Application to industrial case studies showed how the analysis can benefit from the direct inclusion of maintenance into the model. Future work is needed to create a full framework that allows to easily generate maintenance models for FMTs as well as the integration of rare event simulation.

6 ACKNOWLEDGEMENTS

This work has been supported by the STW-ProRail partnership program ExploRail under the project ArRangeer (12238). Further, we like to thank Matijn van Noort from ProRail for the fruitful collaboration.

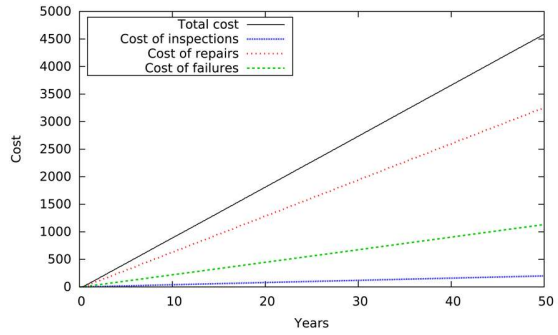

Figure 8 – Breakdown of costs over time for one EI-Joint.

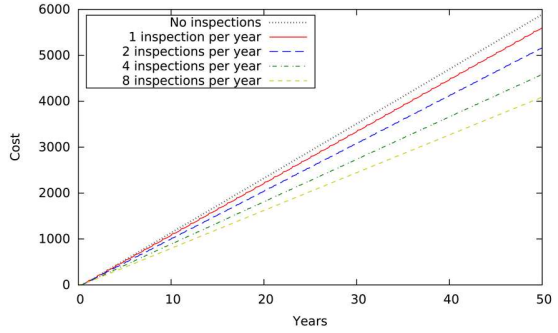

Figure 9 – Effect of different inspection intervals on costs.

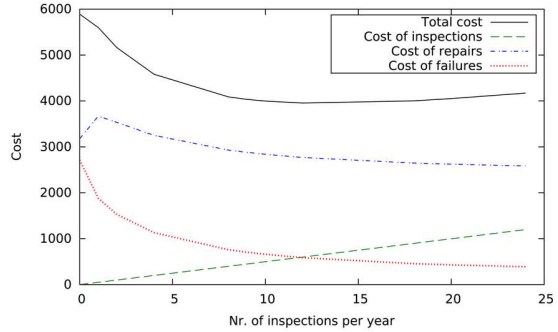

Figure 10 – Effect of different inspection intervals on costs over time.

REFERENCES

1. W. E. Vesely, F. F. Goldberg, N. H. Roberts, and D. F. Haasl, "Fault Tree Handbook". *Office of Nuclear Regulatory Reasearch*, U.S. Nuclear Regulatory Commision, 1981.
2. A. Bobbio and D. Codetta-Raiteri, "Parametric fault trees with dynamic gates and repair boxes," pp. 459–465, *IEEE*, 2004.
3. "Fault tree analysis (FTA)", *Norm IEC 61025:2006(E)*.
4. K. Buchacker, "Modeling with extended fault trees," in *Proc. of the 5th Int. Symp. on High Assurance Systems Engineering (HASE)*, pp. 238–246, Nov 2000.
5. R. Pulungan, "Reduction of acyclic phase-type representations". *PhD thesis*, Universität des Saarlandes, Postfach 151141, 66041 Saarbrücken, 2009.
6. F. Arnold, A. Belinfante, D. G. Freark van der Berg, and M. Stoelinga, "DFTCalc: A tool for efficient fault tree analysis," in *Proc. of the 32nd Int. Conf. on Computer Safety, Reliability and Security (SAFECOMP)*, LNCS, pp. 293–301, Springer, 2013.
7. G. Behrmann, A. David, K. G. Larsen, J. Hakansson, P. Petterson, W. Yi, and M. Hendriks, "Uppaal 4.0," in *Proc. of the 3rd Int. Conf. on Quantitative Evaluation of Systems (QEST),* pp. 125–126, IEEE, 2006.
8. G. Behrmann, K. G. Larsen, and J. I. Rasmussen, "Priced timed automata: Algorithms and applications", *Formal Methods for Components and Objects*, vol. 3657, pp. 162 – 182, 2005.

BIOGRAPHIES

Enno Ruijters
University of Twente, Formal Methods and Tools
Enschede, Overijsel, 7522 NB, The Netherlands

e-mail: e.j.j.ruijters.utwente.nl

Enno Ruijters is a PhD. Student at the University of Twente, currently studying fault tree analysis and stochastic model checking in a railroad infrastructure context. He holds an MSc. in Operations Research from Maastricht University.

Dennis Guck
University of Twente, Formal Methods and Tools
Enschede, Overijsel, 7522 NB, The Netherlands

e-mail: d.guck@utwente.nl

Dennis Guck is a PhD student in the FMT group at the University of Twente. He received his MSc (2012) in computer science from the RWTH Aachen University. His current research includes DFT analysis using model checking.

Peter Drolenga
NedTrain
Utrecht, The Netherlands

e-mail: peter.drolenga@nedtrain.nl

Peter Drolenga will receive his MSc in Applied Physics from the University of Groningen in june 2015. Currently he is working as a researcher at NedTrain at the department Maintenance Development.

Mariëlle Stoelinga
University of Twente, Formal Methods and Tools
Enschede, Overijsel, 7522 NB, The Netherlands

e-mail: marielle@cs.utwente.nl

Dr. Mariëlle Stoelinga is an associate professor at the University of Twente, the Netherlands, where she leads a team on quantitative analysis and risk management of computer systems. She holds an MSc and PhD degree from Radboud University Nijmegen, the Netherlands.