Optimizing fixed-interval maintenance periods for mobile asset sub-systems operating in remote locations

Melinda Hodkiewicz[1], Deion Bonser[2], Alastair Chin[3] and Lyndon While[4]

[1,2,3] *School of Mechanical & Chemical Engineering, The University of Western Australia, Crawley, Western Australia 6009*
melinda.hodkiewicz@uwa.edu.au
c.alastair@hotmail.com
deion.bonser@gmail.com

[4] *School of Computer Science & Software Engineering, The University of Western Australia, Crawley, Western Australia 6009*
lyndon.while@uwa.edu.au

ABSTRACT

Optimizing fixed-interval maintenance periods for mobile asset sub-systems operating in remote locations A typical mobile asset has some number of sub-systems, each with its own maintenance interval and reliability. Our challenge is to set appropriate maintenance intervals for each asset sub-system to minimize the probability of an unplanned failure, without limiting unnecessarily the availability of the asset. To do this we generate sets of maintenance intervals using a genetic algorithm and test them using a discrete event simulation (DES) model of the operations and maintenance functions. We are motivated by two industry examples of fleets in remote locations: long-distance freight trucks, and heavy-haulage rail locomotives. In the truck case, the model found optimal intervals similar to those used by the operator. The locomotive case is more complex, but the model suggests improvements are possible in interval selection, maintenance practices, and data collection. Each model is conceptualized for its specific context; this process identifies assumptions that need to be considered when linking maintenance and operations models.

1. INTRODUCTION

Preventive maintenance is a common strategy used by asset owners to reduce the likelihood of a failure in their assets. Conducting preventive maintenance carries an inherent cost, which we expect to be balanced by an improvement in the asset's reliability. Preventive maintenance is a scheduled endeavor, so a question asked by asset managers who are not bound by warranty considerations imposed by the original equipment manufacturer, is 'What is the optimal interval to conduct preventive maintenance on my asset?'. In this paper we consider the case when mobile assets are deployed in remote locations far removed from the nearest maintenance facility. A failure in-transit results in significant operational delays, additional costs, and unnecessary exposure to hazards for technicians operating in remote and poorly-supported environments. Typical examples of this are trains, road trucks, heavy mobile equipment, cars, and boats. In these situations asset managers seek to reduce the likelihood of unplanned failures of the mobile assets while they are in transit.

Mobile assets comprise a number of functional sub-systems such as engine, drive train, electrical, hydraulic, structural, and so on. These sub-systems are usually represented, from a reliability perspective, as a series system in which the failure of one sub-system causes the failure of the asset. Each sub-system has its own failure modes, associated failure distributions, and planned maintenance (PM) activities and intervals. Modeling of asset fleet availability needs to account for the different maintenance intervals of individual sub-systems of each asset and the stochastic nature of many inputs. For example, the probability and consequence of different failure modes and the asset's travel times all need to be represented by statistical distributions with associated assumptions about uncertainty. In systems with high levels of PM, sub-system life data is often censored so there is limited failure data for analysis meaning that both the actual life and the potential failure modes are at least partly unknown.

Given these uncertainties, our intention is to create a configurable end-to-end optimization process at a level of detail sufficient to deliver optimal PM interval estimates while limiting unnecessary extraneous detail and assumptions. Two case studies explore context-related model-building assumptions and parameter estimation decisions. Impacts of these assumptions on model behaviour are explored and the relation between model and actual system behaviour is examined.

2. LITERATURE REVIEW

Maintenance costs can reach up to 70% of production costs, (Alrabghi & Tiwari, 2015) so improvements in maintenance practices can result in significant economic benefits. This has lead to substantial research on maintenance optimization models. In the early days, approaches focused on analytical models with numerical solvers (Dekker, 1996; Horenbeek et al., 2010; Sharma et al., 2011). While of academic interest, these models did not find wide application for many real-world maintenance problems. Through the 2000s (Rezg et al., 2005; Alrabghi & Tiwari, 2015) reported an increase in use of discrete-event simulation (DES) with genetic algorithms (GAs) being popular for optimization.

GAs are a type of evolutionary algorithm commonly used to solve real-world problems which may be intractable or involve messy, highly non-linear systems (Huband et al., 2003). Evolutionary algorithms are easy to integrate with DES software (Khebbache-Hadji et al., 2012), can search large solution spaces (Marseguerra & Zio, 2000) without getting stuck at local optima (Allaoui & Artiba, 2004; Alrabghi & Tiwari, 2015), can solve multi-objective problems (Moradi et al., 2011), and are often able to solve problems without prior knowledge of the response of the system (Biethahn & Nissen, 1994; Alrabghi & Tiwari, 2015). As a result they have a history of being used successfully to solve challenging problems (Ng et al., 2009; Paulo et al., 2016; Sanchez et al., 2012).

While the number of papers that simulate real case studies has increased, many modelers limit their applications to relatively simple systems. Examples include a single asset (Lhorente et al., 2003), many units of the same asset-type (Bazargan & McGrath, 2003; Berquist et al., 2002; Khebbache-Hadji et al., 2012; Moghaddam, 2013), a single asset with several components (Paulo et al., 2016), or several different assets (Kortelaimen et al., 2000; Louit & Knights, 2001). Nowakowski & Werbinka (2009) explored multi-component systems and their interactions, noting that while this complicates the modelling and optimization process, it also offers the opportunity to group maintenance actions and thus limit the effects of maintenance on system performance. In practice, real-world systems have multiple assets with multiple subsystems, each with different reliabilities, different deterioration profiles, and different failure interactions. There is considerable room for further developments in this area.

In this paper we are interested in mobile assets operating in remote locations. Hence maintenance can only occur at a limited number of places. Failures in transit can take significant time to rectify resulting in delays and associated penalties. Hence, the risk of failure significantly affects the maintenance choices. Maintenance optimization models for these systems will need to be different those more commonly covered in the literature for fixed assets such as in manufacturing (Allaoui & Artiba, 2004) or plant equipment (Berquist et al., 2002).

Aircraft are one example of mobile assets where in-transit failures have financial and possibly catastrophic consequences. As a result, commercial aircraft maintenance strategies are constrained by regulations. For example in the United States, these are set by the Federal Aviation Administration and thus maintenance intervals have limited room for optimization (Sarac et al., 2006). Rail systems are another example of remote assets. Failed locomotives can block other trains from passing and can take significant time and effort to rectify. There has been some research into maintenance optimization of rail systems, but none has assessed the maintenance practices on locomotives specifically. Crainic (2002) surveyed optimization models for long-haul freight cars. Ferreira (1997) highlights the lack of research on delay management, with the primary method of delay prevention being additional investment into tracks, sidings, control systems, and communication systems; but he does not mention adjusting locomotive maintenance practices as a way to reduce the risk of delays.

As Dekker et al. (1997) identified, a lot of research and optimization has been conducted solely to explore the challenges of mathematical analysis, without trying to make the results valuable for real maintenance practitioners. Scarf (1997) recognised this as well, and recommended a shift towards applied research, involving collaboration between researchers and the industry personnel who have the capacity to apply the outcome.

3. APPROACH

The creation of a configurable end-to-end optimization process for estimating optimal PM intervals for a mobile equipment fleet operating in remote locations can be divided into four phases.

Phase 1 involves describing the context of the problem and formulating the fitness function which describes what is to be optimized. In maintenance, commonly used factors in fitness functions are cost, availability, and reliability.

Phase 2 involves the conceptualization of a functional taxonomy for the mobile asset and estimation of relevant maintenance, reliability and operational performance parameters at system and sub-system levels. It is not practical to model every element and operational and maintenance context, so an appropriate level of abstraction is required.

Phase 3 involves the abstraction of the operational context to build a representative simulation of the movement of the asset through the logistics and maintenance systems.

Part 4 involves the development and validation of the genetic algorithm optimization engine, coupled to the simulation.

The decisions involved in each phase are described in more detail below.

3.1. Problem statement and context

The initial stage involves the development of a use case that describes the context of the problem and the objective(s) of the analysis. Common options for the optimization fitness function are cost, availability, or reliability, each of which must be described explicitly. For example, does 'cost' include maintenance activities and the cost of failure? Is 'availability' relevant to the problem of ensuring assets arrive without failure, or should we be interested in 'reliability' for a specific journey. Discussions with the participating organization are necessary to understand the number and configuration of the assets, the logistics and maintenance systems, and how key decisions are made. At this stage an assessment should be made of the quality of the maintenance records and the identification of subject matter experts to support the analysis.

3.2. Conceptualizing the asset

Each mobile asset is modelled as a set of sub-systems, for example sub-systems X, Y and Z. The number and composition of the sub-systems depends on the application. Considerations include the physical and functional structure of the asset, how sub-systems are currently defined and grouped for maintenance, and the reliability of components in the sub-systems. Each subsystem is maintained at a PM interval. These sub-system groupings and associated intervals are set out in the scheduled PM strategy. This strategy may be developed by the equipment operator but is often specified by the Original Equipment Manufacturer (OEM) . For example the OEM might suggest that sub-system X undergoes PM every 100 km, Y every 200 km, and Z every 300 km. Interval units used for mobile assets include distance, calendar time, and utilized time.

Estimation of relevant maintenance, reliability and operational performance parameters at system and sub-system levels requires access to data on the asset, its maintenance schedule, and data from which failure and repair history can be estimated. The preferred approach is to represent actual asset time-to-event and time-to-repair performance as distributions by failure mode at the selected sub-system level. This requires access to historical maintenance records and an ability to interpret the structured and unstructured data they contain. With large data sets this necessitates the use of syntactic analysis skills. In the absence of suitable maintenance data, assumptions can be made based on expert knowledge. The estimated parameters for time-to-event and time-to-repair distributions are usually fit to a Weibull distribution.

Conceptualizing the logistics and maintenance system

We assume mobile assets move along prescribed routes between source(s) and destination(s). This movement is represented using a discrete event simulation model. Typically we require knowledge of a map or schematic of the routes with distances and travel times, factors that might need to be represented such as scheduled stops for refueling and breaks, the location of maintenance shops, maintenance strategies and schedules, and manpower resources.

There are different ways of conceptualizing asset states in the simulation. For example an asset can be assigned one of three states at any given time: 0, completing a journey; 1, in maintenance; or 2, waiting out a rest period. The control flow of the simulation first identifies which asset is closest to its next change of system state by finding the lowest value in the *Distance/Time Remaining* column. All assets are then progressed forward this distance by decreasing their *Distance/Time Remaining* and increasing the age of each of their subsystems. The state of the asset that has reached the state change is then changed depending on a series of checks as shown in Figure 1. If the asset has failed or exceeded its scheduled PM age, maintenance is conducted when it completes its journey. The simulation continues until each asset in the fleet has traveled past a pre-determined test distance.

Alternatively one can consider each asset as having multiple states and moving between states using actions which are scheduled in the simulation environment. Once all assets have scheduled their actions, the environment progresses to the next action to simulate the progression of time. Failures can interrupt this schedule forcing affected assets to re-assess their own state and change their next action. The decision on the number of states and how to transition between them is an important decision for the modeler.

The simulation needs to include provision for PM events at the repair shop and unplanned events (failures) requiring CM in the field. In operating systems there are often undocumented rules that say, for instance, that if one scheduled task is within a certain time of another scheduled task, then they can be done at the same time. These rules need to be elicited and represented in the model. A number of decisions need to be made in developing the simulation such as how to manage component time evolution, how to trigger failure events, how to penalize failure events in the field, and assumptions about the quality of repair. In the case studies described in this paper the virtual age of components is recorded in kilometres in one case, and in hours of operation in the other. The decision of which to use is problem specific. In both cases we schedule a PM task if a subsystem's virtual age has a) passed the planned interval (e.g. traveled 200 km) since the previous planned intervention, or b) if the subsystem's virtual age will exceed the failure age before returning from its next journey and the maintenance workshop is available. The aim of the PM task is to enable the sub-system to function without failure until the next scheduled PM task; however as reliability is stochastic, the sub-system can fail and this is explored through the simulation by sampling from the sub-system's failure distribution. If the sub-system fails during a journey then CM is

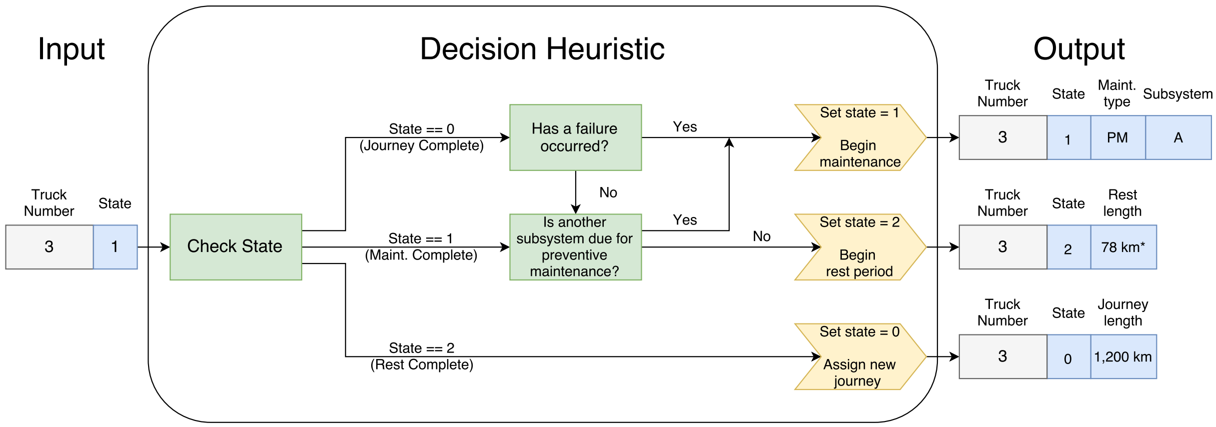

Figure 1. Simulation control flow logic.

conducted. This incurs a cost penalty in the model to capture the issues associated with in-field failures, such as the costs to get maintainers to the remote area, time-delays to receive the product the mobile asset is carrying, fleet schedule disruptions, and the potential for increased safety exposure of the maintainer.

3.3. Optimization, verification and validation

GAs are inspired by the concept of natural selection in biological evolution: they perform an intelligent exploration of a search space of potential solutions by using known good solutions to identify promising new solutions using selection and inheritance. Over some number of iterations, solutions share and accumulate good features, and the performance of the best solutions improves towards the best available. The system terminates usually either when improvement ceases, or when a specific performance level is attained.

The asset fleet is modeled as a multi-objective set of mobile assets which completes a range of journeys under a usage-based PM schedule generated by the GA. The simulation performs trials sending the assets to perform their function executing PM maintenance as per the schedule generated by the GA and dealing with any CM needs that arise.

Outputs from the simulation may include the number of journeys completed successfully, availability and/or cost, and CM or PM conducted per subsystem for each trial schedule. Selected outputs are used by fitness function to calculate a fitness score that represents the success of each schedule. The schedules are ranked and the most successful schedules are sent to the GA to create the next generation of trials as shown in Figures 2 and 3. Over some large number of iterations, the GA will develop new solutions that improve on their predecessors by sharing good features and exploring the search space.

The simulation was developed as a simple state-changing simulation initially before extra functions such as stochastic failure, maintenance checking and the GA were implemented. The GA and simulation are written in Python for one case study and in Java for the other. The first steps to developing the GA are defining how solutions will be represented and how they will be ranked. Each chromosome can be represented as a set of numbers representing the maintenance interval of each subsystem. The fitness of each set of maintenance interval is based on the availability of the system found by the simulation. The crossover and mutation algorithms evolving the population are also defined. Crossover is the process of combining two parent chromosomes to make child chromosomes for the next generation. This is done by choosing a random proportion of each parent maintenance interval and adding them together. Mutation of the children is done by sampling a normal distribution to randomly change the maintenance intervals adding randomness and diversity to the algorithm. Informal static and dynamic VV&T techniques were applied between each development to ensure that the functionality was maintained between updates. A combination of the credibility assessment techniques recommended by Balci (1995) and Kleijnen (1995) were used to verify, validate and test (VV&T) the simulation model. Formal VV&T techniques were not considered viable due to the large confidence intervals of the reliability analysis. With such uncertainty between the true and measured values, comparing the performance of the simulation to that of the real fleet is unable to provide significant support for the credibility of the simulation model. Instead, informal, static, and dynamic (sensitivity analysis) VV&T techniques are used to support the credibility of the simulation. Informal techniques are common in VV&T practices (Balci 1990). While they rely more on human judgment rather than mathematical analysis (Balci 1997), they still have a structure and can be effective.

4. RESULTS AND DISCUSSION

Two case studies are provided. The first is on five long distance trucks. These trucks haul supplies between a city and

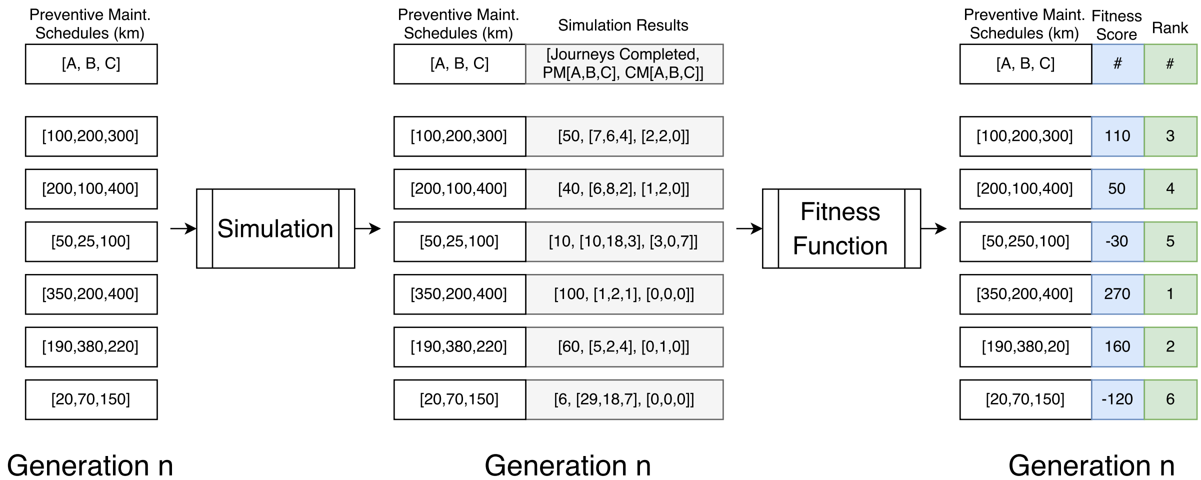

Figure 2. Simulation and optimization process.

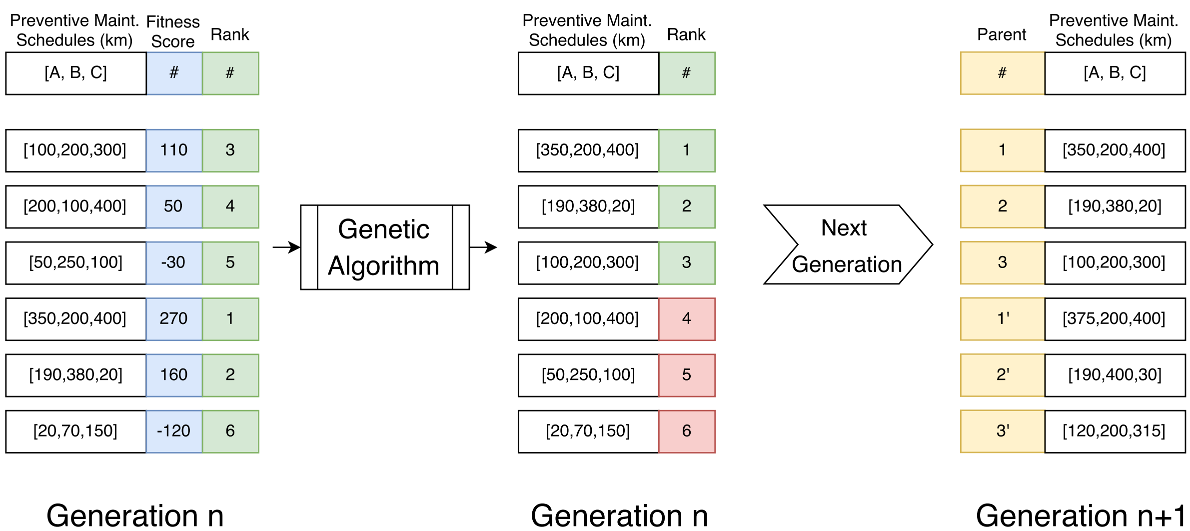

Figure 3. Genetic algorithm process.

a set of ten remote towns, which are between 390 km and 1400 km away. The second is for the locomotives on a heavy haulage railway, the travel time between source and destination is between 32 and 45 hours. In both cases there is almost no maintenance support between the source and destination and the environment is hostile in that temperatures frequently exceed 40°C and there are no power or mobile network services and few towns.

4.1. Long distance trucks

4.1.1. Model development

A company operates five long-distance haulage trucks, also known as prime-movers or road trains. The aim is to minimize the sum of the PM and CM costs through setting appropriate intervals for PM on the truck sub-systems. The costs of PM and CM events for each subsystem are assumed for the purposes of this model to be fixed values, though in practice there are a number of factors that can influence these costs. The number of PM and CM events are determined by the simulation and the intervals for the PM determined by the GA.

The simulation model simulates the operation of the trucks and their associated subsystems. It includes a model for different journey options and a single maintenance workshop. The workshop is treated as the central node from which each truck begins and ends journeys. It is only capable of conducting maintenance on one truck at a time. A good-as-new maintenance quality model was used for the baseline simulation. In this model, the maintenance quality is considered to be 'perfect' and each PM and CM action returned the age of the subsystem to a 'good-as-new' state.

A taxonomy of the trucks was developed in consultation with experts and the current maintenance practices documented.

Component Name	Subsystem	β	η (km)
Pneumatic System	A	1.36	65,681
Hydraulic System	A	2.22	69,058
Lights	A	1.11	71,618
Hoses	A	1.28	99,513
Sensors	A	1.37	103,399
Tyres	B	1.12	138,047
Cabin	B	1.31	150,691
Wiring	B	1.65	181,044
Exhaust	B	1.67	281,213
Turbocharger	C	1.33	640,924
Shaft	C	1.07	792,508
Engine	C	1.28	1,146,408
Radiator	C	1.05	2,396,505

Table 1. Failure probability Weibull parameter estimates for truck components.

Service reports listed what PM had been conducted, what components had failed and whether they were repaired or replaced. Historical service reports for an 18-month period were available in two formats (Excel spreadsheet and .pdf files). These files were collated in a single dataset listing the time-to-event (in kilometres) and parts replaced for each maintenance (planned or failure) event. Only a small number of failures were observed in the data due to the the short time on test. When a component was repaired or replaced, it was treated as a failure event that caused the truck to break down in the field. When the beginning or end of life did not occur during in the 18-month test window this censoring was reflected in the estimation of parameters for the reliability distribution. Reliability analysis was performed in R using the open source code at (Marriott (2016)) based on the method proposed by (Meeker & Escobar (1998)). This analysis provided lifetime and shape parameter estimates for a two-parameter Weibull distribution to represent the reliability of each component. Three sub-systems (A, B and C) were created by grouping the components based on similar Weibull scale parameter values (η). The resulting reliability block diagram in a series configuration is shown in Figure 4. The subsystems produced by this approach are different to the maintenance subsystem grouping used by the owner of the prime mover fleet based on recommendations of the original equipment manufacturer.

The reliability data used in the model is shown in Table 1. The simulation begins by initializing each vehicle subsystem with a random age up to its projected failure age. Each mobile asset is assigned a journey randomly selected from a predefined set of journeys based off the fleet's real operation and a new maintenance schedule generated by the optimization engine described below. A two-dimensional array records the progress and state of each prime mover in the fleet throughout the simulation (Table 2). The simulation is distance-based (km); time-based events such as maintenance are converted to distances using an estimate for average truck speed.

4.1.2. Results

The results from the simulation-optimization model for each sub-system A, B and C are shown in Figure 5. For each subsystem there is a distribution of high-scoring schedules. A high-scoring schedule has maximized the fitness score. From these results a suitable maintenance schedule is to repair subsystem A every 12,000 km, B every 34,000 km, and C every 180,000 km. An additional consideration is the potential to extend the maintenance interval for subsystem B to 36,000 km so that it lines up with every third time maintenance A is conducted. Maintenance on subsystem C would then be conducted along with A and B every 15^{th} time maintenance A is conducted.

To explore the behaviour of the model we reduced a) the lengths of journeys, b) the cost value of a failure, and c) failure duration. This would be the case if the trucks were operated in a less-remote environment. Five trials were conducted using this adjusted simulation, which all converged to a similar range of optimal schedules to the original result. Across these trials,the optimal preventive maintenance schedule identified by the GA increased slightly. Both the changes to the cost of a failure and the shorter journey length are expected to be contributors to this phenomena. Since the cost of a failure has been decreased relative to the cost of conducting preventive maintenance, the GA began favouring schedules that allows the subsystems to reach a slightly higher age before being repaired. The shorter maintenance duration due to a reduced maintenance response time also allows the vehicle to begin conducting journeys again for profit sooner than in the original simulation.

The change to conducting shorter journeys was also expected to cause a small increase in the optimal maintenance schedules. For example, if a truck in the original simulation was due for preventive maintenance in 1,000 km and would need to conduct a 1,300 km journey, it would conduct the maintenance at that age before completing the journey. However, if that same truck was due for maintenance in 1,000 km in this adjusted simulation with shorter journey's, it would still be able to conduct many shorter (e.g. 20–40 km) journeys before being due for maintenance.

The PM intervals suggested by the truck model compare favorably to the current strategy used in the real fleet where maintenance A, B, and C is conducted every 10,000 km, 20,000 km, and 100,000 km respectively. It is interesting to note that although the sub-systems were grouped into A, B, and C based on the mean time between failure determined from maintenance data and expert that the intervals suggested by this analysis are very close to those recommended by the truck's Original Equipment Manufacturer. We would expect some differences in interval predictions due to the assumptions and simplifications in the model. The value of the model lies in the suggestion that maintenance intervals could be ex-

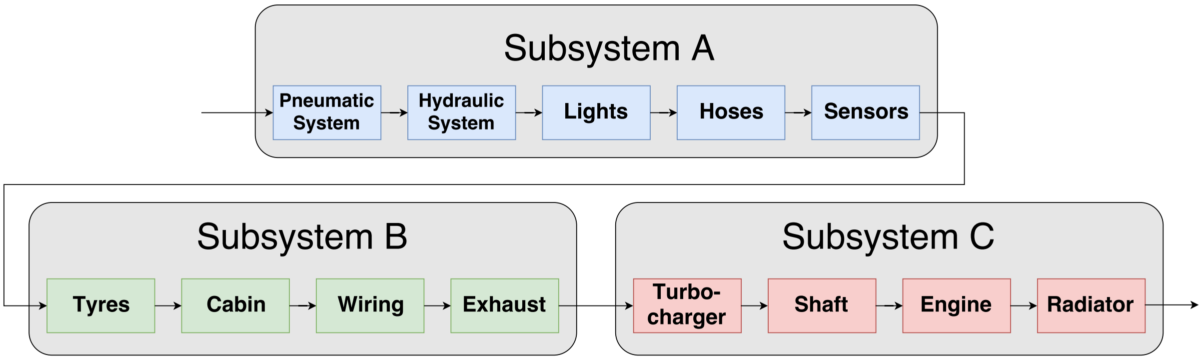

Figure 4. Subsystem grouping based on reliability analysis.

Truck ID	Current Sys State	Dist./Time Rem. (km)	Subsystem A		Subsystem B		Subsystem C	
			Age	Fail	Age	Fail	Age	Fail
1	0	1,234	17,000	30,574	89,375	160,734	178,263	320,593
2	1	39	3,882	23,718	20,407	124,693	40,703	248,707
3	0	688	6,862	23,703	74,301	192,613	49,492	169,394
4	0	184	7,798	26,403	40,996	138,805	81,769	276,854
5	2	322	2,426	22,573	12,752	118,671	25,435	236,696

Table 2. Representation of instantaneous truck fleet data in the simulation at different states: 0, completing a journey; 1, in maintenance; or 2, waiting out a rest period.

tended and the opportunity, through the model, to identify which areas to improve on and the scale of the improvement that could be expected.

4.2. Heavy haulage locomotive

4.2.1. Model development

This asset system contains 118 locomotives operating in a train network. Each train comprises four locomotives and more than 100 wagons. The aim of the project is to maximize the availability of the system through setting appropriate intervals for PM for sub-systems on each of the 118 locomotives. The same PM interval is used for each identical sub-system. The reliability of the wagons is not included.

In this case study we adopted the sub-system structure provided by the company. Each locomotive is modelled as six subsystems: engine, electrical (low voltage), electrical (high voltage), drive, pneumatics, and structure. Data about the locomotives was received in the form of two years worth of maintenance work orders. These 21,076 work-orders contain a description of the work, the time of occurrence, how long the action took, what component/sub-system of the locomotive was maintained, and how much each action cost. The data needed considerable cleaning using syntactic analysis to identify the component/sub-system, date of maintenance activity, type of activity, age of component at removal and if the end of life constituted a failure or a suspension. Parameters for the failure distributions for the locomotive sub-systems are shown in Table 3.

The times to conduct PM and CM on the subsystems are determined from the historical work order data and their distributions shown in Figures 6 and 7. In addition to data for the model, these graphs shows which subsystems take the longest time to maintain. PM repair time distributions are represented by normal or log normal distributions while CM repair time distributions usually need a mixed Weibull fit. On average, a corrective repair on a locomotive subsystem takes more than six times as long as planned repair. This highlights the cost of failures to assets and reinforces why maintenance optimization is important to minimize the unplanned failure of assets.

The model simulates the movement of the 118 trains (each requiring four locomotives) traveling between several sources and a single destination. The maintenance shop is at the destination. Cycle times are assumed to follow a triangular distribution. Failures on the sub-systems are assumed to be independent of other sub-systems and a sub-system is considered as good-as-new after a repair event. The simulation control logic is quite complicated because of multiple sources, single tracks with limited passing options and the number of locomotives involved. The fitness function of availability is determined as the useful simulation time the locomotives were operating normalized by the total simulation time. The simulation is run for 50,000 hours, or about six years. This allows maintenance to occur at least twice and still allows a simulation to execute quite quickly, although multiple simulations can take days. Verification of the modelling was done through sensitivity analysis. Tests completed included increasing the

Subsystem	Distribution	Population proportion	Shape	Scale	MTTF (hrs)
Engine	Mixed	0.58	0.84	858	938
	Weibull	0.42	0.67	1.6×10^{13}	2×10^{13}
Electrical (low voltage)	Mixed	0.52	0.92	666	691
	Weibull	0.48	0.79	30,703	35,126
Drive	Lognormal	1.00	2.52	7	34,803
Electrical (high voltage)	Lognormal	1.00	2.30	7	15,321
Pneumatics	Mixed	0.61	0.86	716	774
	Weibull	0.39	0.37	1.32×10^{59}	5.65×10^{59}
Structure	Lognormal	1.00	2.79	8	216,710

Table 3. Calculated failure distributions of locomotive subsystems.

	Subsystem	Existing	GA
Maintenance Interval (hrs)	Engine	4,842	34,866
	Low voltage	7,635	38,479
	Drive	8,471	29,261
	High voltage	8,166	21,017
	Structure	5,304	35,292
	Pneumatic	8,177	34,558
	Availability	35.5%	62.0%

Table 4. GA-determined maintenance intervals for each maintenance scheduling algorithm.

number of locomotives from 20 to 200, and observing the impact on system availability, locomotive idle time, and the number of trips completed. In each case we look at the behaviour of the system. For example, as the number of locomotives is increased, Idle Time decreases initially as the system utilizes these extra assets; it then increases slightly as there are not enough trains to utilize the number of locomotives; and finally it decreases further as the queues for maintenance and repair become very long. All of this accords with our intuition of the system.

4.2.2. Results

A baseline for the model is established by running the GA over maintenance intervals between 6 and 12 months reflecting the existing maintenance program and also up to 36 months with the option for the GA to mutate intervals above this value. Using the GA to set intervals for each of the six subsystems as shown in Table 4 results in a 75% better availability when compared to the baseline model.

The model suggests a run-to-failure strategy because the shape factor on the reliability distributions for most of the subsystems (determined from the work order data) are less than 1. As a result the model does not try to reduce the risk of in-transit failures. Since the data suggest the subsystems suffer from infant mortality the model allows them to fail as this affects the availability fitness function less than the maintenance required to prevent these in-transit failures. A different fitness function, e.g. one based on cost, would produce a different result; this will be explored in future work.

The model's results suggest that maintenance intervals for the locomotives can be extended, as the simulations in Table 4 suggest that availability is higher with longer maintenance intervals. However more work is required to confirm the data on which the model is built. There were no trustworthy parameters concerning the reliability distributions of the subsystems used in the model. We used data from the work order system which required considerable syntactic analysis by the researchers. A shape factor of less than one for some of the reliability distributions is an unwelcome result, Table 3, and is now being investigated by the organization. The organization had not previously calculated reliability parameters for the sub-systems taking suspensions into account and there was no ground truth available on the actual locomotive availability, unplanned downtime or time to repair or maintain estimations. We also suspect that some of sub-systems have maintenance intervals of more than two years. This was difficult to determine given that we had only two years of work order data.

While the locomotive system analysed conducts regular scheduled maintenance, condition monitoring of the locomotive subsystems also takes place. This is done through inspections during stops, during maintenance of other subsystems and by observations of the operators. This means failures can sometimes be pre-empted and parts can be replaced before they fail. The effect of this is less failures and thus, increased availability and production. The effect of the frequency of inspections would have an influence over the optimal PM scheduling. However, the simulation does not take this into account. Optimization done using a combination of both maintenance strategies would further reflect the actual operation of the real system.

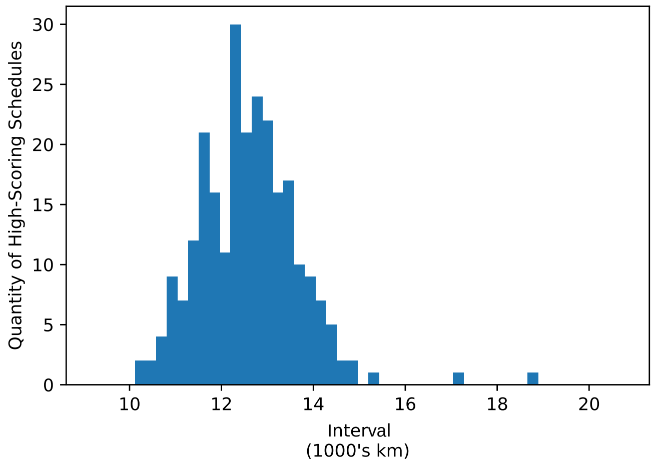

(a) Subsystem A.

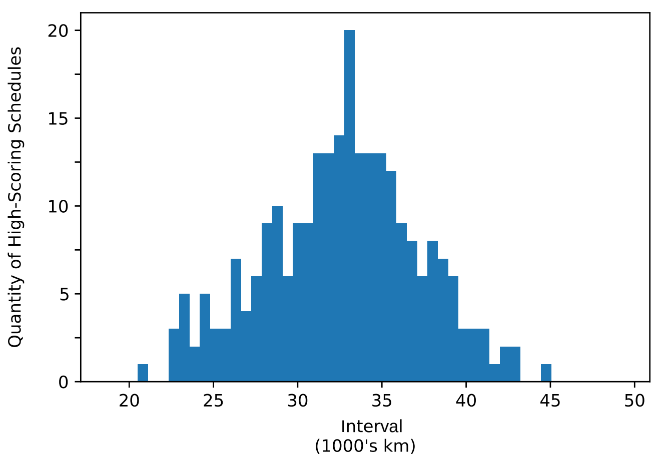

(b) Subsystem B.

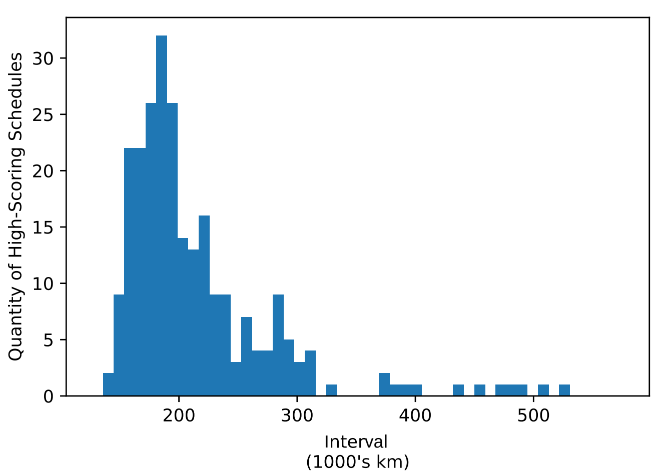

(c) Subsystem C.

Figure 5. Top performing maintenance intervals in the truck simulation

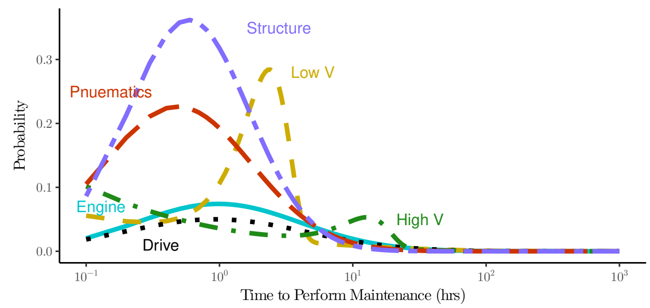

Figure 6. Comparison of probability distribution functions of PM repair times of the locomotive subsystems

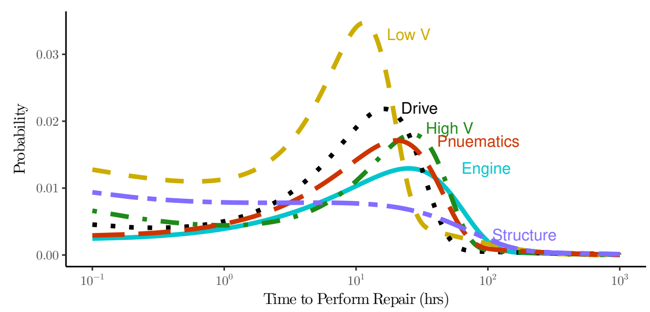

Figure 7. Comparison of probability distribution functions of CM repair times of the locomotive subsystems

4.3. Discussion

The approach has been applied to two cases in different business contexts and on different scales. The heavy-haulage truck case is a relatively simple five truck, three sub-systems per truck case, with very little actual maintenance work order data available and a high reliance of expert knowledge. There is little complexity in the logistics model. The results of the simulation-optimization compare well with actual practice.

The locomotive case is much more complex, 118 locomotives each with six sub-systems, and code which takes 3-4 days to run. This case is much less easy to validate because the logistics simulation is an abstraction of the real, and more complex, rail network, and because there is no ground truth from the operator on the actual reliability or availability of the locomotives. However models like this that highlight where problems exist and provide a mechanism to quantify the value of improved infrastructure and data collection to support this sort of modeling for decision support.

Maintenance and operational data necessary for building the simulation came as extracts of .csv fies from internal corporate systems and discussions with experts. Open source software was used for both the cases. For the long-distance truck case the simulation and GA were developed in Python, and reliability analysis was done in R. For the locomotive case study the maintenance data processing and syntactic analysis was done in Python, and the simulation and GA were written in Java. The point being made here is that building and maintaining these models usually requires capability with different software; in our case Python, R, Java, Weibull++ were all used, as well as an ability to work with industry experts and their systems.

5. CONCLUSIONS

This paper has presented an approach based on coupled simulation and GA models to identify fixed-interval maintenance periods for mobile assets operating in remote locations. The approach has been demonstrated on two case studies, one relatively simple case of five trucks each with three sub-systems, and the other more complex case of 118 locomotives each with six sub-systems. Maintenance records are used to develop distributions for the time-to-event and time-to-repair variables. Both models suggest improvement in fleet availability are possible with alternate interval selection.

6. ACKNOWLEDGEMENTS

The authors would like to acknowledge the support of staff working for the two asset owners who provided data and motivation to support the development of these models. This work was supported in part by the BHP Fellowship for Engineering for Remote Operations

REFERENCES

Allaoui, H., & Artiba, A. (2004). Integrating simulation and optimization to schedule a hybrid flow shop with maintenance constraints. *Computers and Industrial Engineering*, *47*(4), 431 - 450.

Alrabghi, A., & Tiwari, A. (2015). State of the art in simulation-based optimisation for maintenance systems. *Computers and Industrial Engineering*, *82*, 167 - 182.

Balci, O. (1990). *'Guidelines for successful simulation studies', Proceedings of the 22nd Conference on Winter Simulation, pp. 25-32.* Available from: http://informs-sim.org/. [2 October 2017].

Balci, O. (1995). *'Principles and techniques of simulation validation, verification, and testing', Proceedings of the 1995 Winter Simulation Conference, pp. 147-154.* Available from: http://informs-sim.org/. [2 October 2017].

Balci, O. (1997). *'Verification, validation and accreditation of simulation models', Proceedings of the 1997 Winter Simulation Conference, pp. 135-141.* Available from: http://informs-sim.org/. [2 October 2017].

Bazargan, M., & McGrath, R. N. (2003). Discrete event simulation to improve aircraft availability and maintainability. *Annual Reliability and Maintainability Symposium.*

Berquist, S., Buckley, A., & Bodon, P. (2002). Using simulation modeling to evaluate maintenance changes. *ICOMS-2002*, 1-11.

Biethahn, J., & Nissen, V. (1994). Combinations of simulation and evolutionary algorithms in management science and economics. *Annals of Operations Research*, *52*(4), 181–208.

Crainic, T. G. (2002). *A survey of optimization models for long-haul freight transportation* (2nd ed.). QC, Canada: Kluwer.

Dekker, R. (1996). Applications of maintenance optimization models: a review and analysis. *Reliability Engineering and System Safety*, *51*(3), 229-240.

Dekker, R., Wildeman, R. E., & van der Duyn Schouten, F. (1997). A review of multi-component maintenance models with economic dependence. *Mathematical Methods of Operations Research*, *45*(3), 411-435.

Ferreira, L. (1997). Planning Australian freight rail operations: An overview. *Transportation Research Part A: Policy and Practice*, *31*(4), 335 - 348.

Horenbeek, A. V., Pintelon, L., & Muchiri, P. (2010). Maintenance optimization models and criteria. *International Journal of Systems Assurance Engineering*, *1*(3), 189-200.

Huband, S., Hingston, P., While, L., & Barone, L. (2003). An evolution strategy with probabilistic mutation for multi-objective optimisation. In *2003 congress on evolutionary computation* (Vol. 4, p. 2284-2291).

Khebbache-Hadji, S., Hani, Y., Lahiani, N., & Mhamedi, A. (2012). Genetic algorithm used in simulation model: Application in a maintenance process. *{IFAC} Proceedings Volumes*, *45*(6), 1047 - 1052. (14th {IFAC} Symposium on Information Control Problems in Manufacturing)

Kleijnen, J. (1995, 02). Verification and validation of simulation models. *European Journal of Operational Research*, *82*, 145-162.

Kortelaimen, H., Salmikuukka, J., & Puriso, S. (2000). Reliability simulation model for systems with multiple intermediate storages. *Annual Reliability and Maintainability Symposium*.

Lhorente, B., Lugtigheid, D., Knights, P. F., & Santana, A. (2003). A model for optimal armature maintenance in electric haul truck wheel motors: a case study. *Reliability Engineering & System Safety*, *84*, 209-218.

Louit, D. M., & Knights, P. F. (2001). Simulation of initiatives to improve mine maintenance. *Mining Technology*, *110*(1), 47,58.

Marriott, R. (2016). *Reliability analysis in R.* Available from: http://uwa.engineering/project/example-of-a-reliability-analysis-in-r/. [24 September 2017].

Marseguerra, M., & Zio, E. (2000). Optimizing maintenance and repair policies via a combination of genetic algorithms and Monte Carlo simulation. *Reliability Engineering and System Safety*, *68*(1), 69-83.

Meeker, W. Q., & Escobar, L. A. (1998). *Statistical methods for reliability data* (1st ed.). New York: John Wiley and Sons, Inc.

Moghaddam, K. S. (2013). Multi-objective preventive maintenance and replacement scheduling in a manufacturing system using goal programming. *International Journal of Production Economics*, *146*(2), 704 - 716.

Moradi, E., Ghomi, S. F., & Zandieh, M. (2011). Bi-objective optimization research on integrated fixed time interval preventive maintenance and production for scheduling flexible job-shop problem. *Expert Systems with Applications*, *38*(6), 7169 - 7178.

Ng, M., Lin, D., & Waller, S. (2009). Optimal long-term infrastructure maintenance planning accounting for traffic dynamics. *Computer-Aided Civil and Infrastructure Engineering*, *24*(7), 459 - 469.

Nowakowski, T., & Werbinka, S. (2009). On problems of multicomponent system maintenance modelling. *International Journal of Automation and Computing*, *6*(4), 364.

Paulo, P., Branco, F., de Brito, J., & Silva, A. (2016). BuildingsLife – the use of genetic algorithms for maintenance plan optimization. *Journal of Cleaner Production*, *121*, 84 - 98.

Rezg, N., Chelbi, A., & Xie, X. (2005). Modeling and optimizing a joint inventory control and preventive maintenance strategy for a randomly failing production unit: Analytical and simulation approaches. *International Journal of Computer Integrated Manufacturing*, *18*(2), 225-235.

Sanchez, E., Squillero, G., & Tonda, A. (2012). *Industrial applications of evolutionary algorithms*. Springer-Verlag.

Sarac, A., Batta, R., & Rump, C. (2006). A branch-and-price approach for operational aircraft maintenance routing. *European Journal of Operational Research*, *175*(3), 1850 - 1869.

Scarf, P. A. (1997). On the application of mathematical models in maintenance. *European Journal of Operational Research*, *99*(1), 493-506.

Sharma, A., Yadava, G. S., & Deshmukh, S. G. (2011). A literature review and future perspectives on maintenance optimization. *Journal of Quality in Maintenance Engineering*, *17*(1), 5-25.