

Optimizing a Mine Haul Truck Wheel Motors' Condition Monitoring Program: Use of Proportional Hazards Modeling

A.K.S. Jardine*, D. Banjevic*, M. Wiseman[§], S. Buck[†], T. Joseph*

**Condition-Based Maintenance Laboratory, Department of Mechanical and Industrial Engineering, University of Toronto*

§PricewaterhouseCoopers, Toronto

†Cardinal River Coals Ltd., Alberta

Abstract:

The paper discusses work completed at Cardinal River Coals in Canada to improve the existing oil analysis condition-monitoring program being undertaken for wheel motors.

Oil analysis results from a fleet of 55 haul truck wheel motors were analyzed along with their respective failures and repairs over a nine-year period. Detailed data cleaning procedures were applied to prepare data for modeling. In addition, definitions of failure and suspension were clarified depending on equipment condition at replacement. Using the proportional-hazards model (PHM) approach, the key condition variables relating to failures were found from among the 19 elements monitored, plus sediment and viscosity. Those key variables were then incorporated into a decision model that provided an unambiguous and optimal recommendation on whether to continue operating a wheel motor or to remove it for overhaul on the basis of data obtained from an oil sample.

Wheel motor failure implied extensive planetary gear or sun gear damage necessitating the replacement of one or more major internal components in a general overhaul. The decision model, when triggered by incoming data, provided both a recommendation based on an optimal decision policy as well as an estimate of the unit's remaining useful life (RUL). By optimizing the times of repair as a function both of age and condition data a 20-30% potential savings in overhaul costs over existing practice was identified.

Keywords: Wheel motors, Condition monitoring, Oil analysis, Proportional-hazards modeling, Optimizing condition-based maintenance decisions, EXAKT software

Practical Implications: Current practice for monitoring the health of items is through examining trends in readings obtained from various forms of condition monitoring. Interpretation of these readings is undertaken by an inspector reviewing current and past readings, or through using commercially available trending software. Such an approach does not guarantee that the full information-value contained in the readings is captured. The paper uses a statistical procedure called proportional-hazards modeling to identify the key measurements that should be used to assess the true state of health of the equipment. Economic decision rules are then established. The procedure is described

through a case study that reports on the optimization of condition-based maintenance decisions for haul truck wheel motors that are monitored through oil analysis. Application of the procedure demonstrated a 20-30% potential savings in overhaul costs compared to current practice.

This paper underlines the importance of data cleaning and applying a consistent definition of failure based on both the observed equipment condition at repair time and the inability of the equipment to perform its functions. (for additional discussion see Campbell and Jardine 2001).

INTRODUCTION

Cardinal River Coals Ltd. is a 50/50 joint venture between Luscar Ltd. and Consol of Canada, Inc. The mine is located approximately 50 km south of Hinton, Alberta on the eastern slopes of the Rocky Mountains. The coal produced from the mine is low sulphur, high quality coking coal used for steel making. Cardinal River Coals Ltd. opened in 1970 as a multiple open pit mine using the truck and shovel mining method. Current annual production at the mine calls for the removal of 21 million cubic metres of rock and 2.8 million tonnes of coal. The mine has won multiple awards for the land reclamation and creating wildlife habitat.

There are 26 haul trucks at the mine site, each having two wheel motors. With 3 spare wheel motors the fleet numbers 55. The existing policy, based on experience, is to rebuild the units after about 20,000 hours of operation. Oil analysis is carried out monthly whereby the amount of sediment (weight of filter patch filtrate) and parts per million (ppm) of five out of the nineteen elements are noted: iron, silicon, chrome, nickel, and titanium. The decision to remove the unit for rebuild is based on manual perusal of the values of these elements in combination with the unit's age.

Wheel motor failures relating to the electrical drive elements and breaking system were not included in this study since their condition is not reflected by oil analysis data. Seal replacements were carried out frequently as a result of high contamination and coincided with oil changeouts. The oil changeout event (OC) is considered as a "minor" repair. The analysis shows that a high amount of sediment persisting in spite of these corrective measures, is associated with a high risk of failure.

Statistical analysis of the CRC wheel-motor data showed a high correlation between iron and silicon. That fact would support the view that there are a high number of failures which are contaminant induced. Hence one may conclude that there is an event or set of conditions that initiate a process of deterioration in the wheel motor. It is assumed that by overhauling the unit before the damage becomes more extensive one would benefit from savings through failure avoidance.

DATA AVAILABILITY

Within the mine's computer maintenance management system (CMMS) there were histories of wheel motor lifetimes, including details of removals due to failure or preventive maintenance as a result of interpretation of the signals obtained from oil analysis. Costs associated with the failure and preventive removals were also available. Additionally, there was a database containing a vast history of condition monitoring test results – some 50,000 records.

It may seem that it would be an easy matter to peruse and study these two data sources and learn which patterns of data have been associated with past failure, thus identifying the data combinations that might be employed as condition indicators of future failures. Unfortunately identification of the key condition indicators from amongst all the data collected is seldom obvious to the analyst. The complexity, volume, and time lags within the data render them elusive if not impossible to discern without the proper tools.

In this paper we show a tool that uses a statistical modeling technique known as proportional-hazards modeling to bridge these two invaluable data sources. It is the central function in a program called EXAKT developed precisely for this purpose by the condition based maintenance (CBM) laboratory at the University of Toronto (see Jardine et al, 1997).

MODEL BUILDING

The Proportional-hazards Model

A valuable statistical procedure to estimate the risk of equipment failing when it is subject to condition monitoring is the proportional hazards model (PHM)(Cox, 1972). The form of the PHM is:

$$h(t) = \frac{\mathbf{b}}{\mathbf{h}} \left(\frac{t}{\mathbf{h}} \right)^{b-1} e^{g_1 Z_1(t) + g_2 Z_2(t) + \dots + g_m Z_m(t)} \quad (1)$$

where $h(t)$ is the (instantaneous) conditional probability of failure at time t , also known as the hazard function, given the values of $Z_1(t), Z_2(t), \dots, Z_m(t)$. Each $Z_i(t)$ in equation (1) represents a monitored condition data item at the time of inspection, t , such as the parts per million of iron or the vibration amplitude at the second harmonic of shaft rotation. These condition data are called covariates.

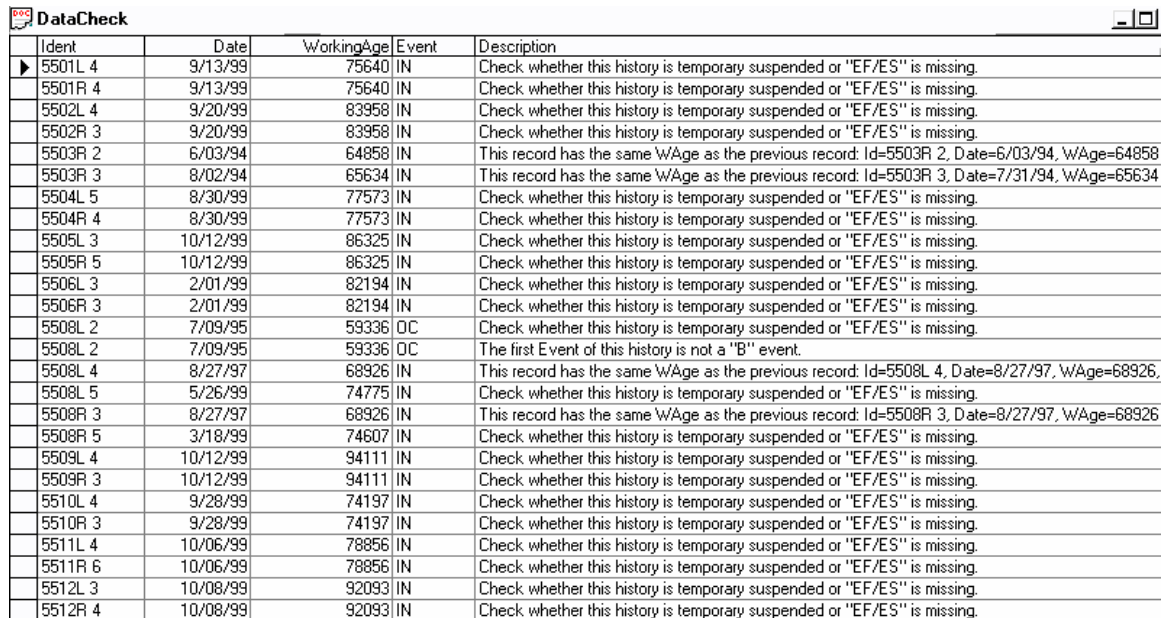
The γ 's are the covariate parameters indicating the degree of influence each covariate has on the hazard function. The model consists of two parts, the first part is a baseline hazard function that takes into account the age of the equipment at time of inspection, $\frac{\mathbf{b}}{\mathbf{h}} \left(\frac{t}{\mathbf{h}} \right)^{b-1}$,

and the second part, $e^{g_1 Z_1(t) + g_2 Z_2(t) + \dots + g_m Z_m(t)}$, takes into account the key variables and their associated weights.

Data Cleaning Related to Wheel Motor Event Records

The first step in every proportional hazard model (PHM) building exercise is a thorough examination of the data. The data-cleaning phase of PHM is considered the most important one of the entire modeling process. If we are to accomplish our objective of accurate and automated CBM data interpretation, the data upon which we intend to build the model must be as free of error as is possible. Much of this paper, therefore, focuses on the data investigative or cleansing phase of model building.

Fortunately software provides us with ample tools with which to “cleanse” the data. In Figure 1 we show a feature of the EXAKT software (2) that discovers many logical data inconsistencies emanating from the CMMS, thereby helping the analyst to make the corrections that will improve the ultimate model’s precision.



Ident	Date	WorkingAge	Event	Description
5501L 4	9/13/99	75640	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5501R 4	9/13/99	75640	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5502L 4	9/20/99	83958	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5502R 3	9/20/99	83958	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5503R 2	6/03/94	64858	IN	This record has the same WAge as the previous record: Id=5503R 2, Date=6/03/94, WAge=64858
5503R 3	8/02/94	65634	IN	This record has the same WAge as the previous record: Id=5503R 3, Date=7/31/94, WAge=65634
5504L 5	8/30/99	77573	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5504R 4	8/30/99	77573	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5505L 3	10/12/99	86325	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5505R 5	10/12/99	86325	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5506L 3	2/01/99	82194	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5506R 3	2/01/99	82194	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5508L 2	7/09/95	59336	OC	Check whether this history is temporary suspended or "EF/ES" is missing.
5508L 2	7/09/95	59336	OC	The first Event of this history is not a "B" event.
5508L 4	8/27/97	68926	IN	This record has the same WAge as the previous record: Id=5508L 4, Date=8/27/97, WAge=68926.
5508L 5	5/26/99	74775	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5508R 3	8/27/97	68926	IN	This record has the same WAge as the previous record: Id=5508R 3, Date=8/27/97, WAge=68926
5508R 5	3/18/99	74607	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5509L 4	10/12/99	94111	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5509R 3	10/12/99	94111	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5510L 4	9/28/99	74197	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5510R 3	9/28/99	74197	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5511L 4	10/06/99	78856	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5511R 6	10/06/99	78856	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5512L 3	10/08/99	92093	IN	Check whether this history is temporary suspended or "EF/ES" is missing.
5512R 4	10/08/99	92093	IN	Check whether this history is temporary suspended or "EF/ES" is missing.

Take in Figure 1 Data Checking Tool

Data required for PHM analysis consists of "histories". Each valid history for a wheel motor must have a Beginning event (B), an Ending event (EF for failure, or ES for suspension (such as a preventive removal)) and Inspection events. A discussion of how suspensions and failures were determined is given later in this paper. A history could also have events that are known to affect covariates, such as oil change (OC) events.

The output of the data - checking tool (illustrated on Figure 1) points out probable errors based on a systematic evaluation of working ages and corresponding calendar dates as reported in the CMMS. Thus it is seen that the software deduces, from the dates and working ages, the sets of data that comprise individual histories. For each history that it finds without an ending, it asks whether the ending event should be designated as a


suspension (ES), a temporary suspension (TS, which is denoted by *ES in the software) or a failure (EF). (Note: Temporary suspension means the age of the operating item at the time of the data analysis. In the future, TS will convert to an ES or EF event.)

The software also points out anomalies that may indicate data problems such as two inspections on the same day, or working ages and calendar dates which are out of synchronization relative to the previous and next records. All of these errors would have compromised the model's accuracy.

Most of these types of errors can easily be corrected by inserting the missing Beginning and Ending events for each history.

Data Cleaning Related to Wheel Motor Condition Monitoring Records

Examination of the records obtained at oil analysis can be examined graphically in many ways using various combinations of covariates, dates, ranges, and scales. For example, while exploring the covariates, statistically unusual values of silicon forming a horizontal line at exactly 900 parts per million (PPM) were noted (Figure 2).




Take in Figure 2 Graphical Analysis of Inspection Records

Investigating with the commercial laboratory it was found that for a period of time the photo-multiplier tube on the spectrometer was saturating at exactly 900 PPM. In other words all values of silicon above 900 were truncated to 900 PPM. A similar situation occurred for iron above 2500 PPM. If not detected, this could play havoc with the building of the PHM.

Knowing the errors in the laboratory test data it was possible to compensate for them in the database used to build the model. For example, to correct the truncated values of 'Si' they were replaced with $1.2 \times \text{Fe}$. The factor of 1.2 was determined from the initial slope of the cross graph (a correlation graph) of Fe-Si and values obtained after the saturation

defect was corrected. The truncated Fe values were not corrected since there were too few of them to influence the model.


The correction applied to the Si values is illustrated in Figure 3.



Cross graphs of pairs of covariates are invaluable in finding correlations that are of great help in developing and evaluating the eventual model.

Figure 4 shows the correlation between iron (Fe) and nickel (Ni). Correlations between other covariates were also tested. For example, Fe vs Ti, Fe vs Si and Ni vs Ti graphs all exhibited similar correlation.


Determining correlation between covariates is useful both to provide insight into the data, and in understanding the models generated by the software. For example, if 'Fe' and "Ni" are highly correlated the modeling process would indicate that there is no point in including nickel in the model since it has been determined to provide no additional information regarding the probability of failure. Thus, if the software concludes that nickel is "insignificant", then by inspecting the correlation graphs one could therefore understand the reasonableness of such an indication. These correlations are the result of wear of a metallic alloy component present in the unit.



Take in Figure 4 Correlating Iron and Nickel

Data Cleaning Related to Wheel Motor Oil Changes

When building the PHM it is necessary that when the inspection records are analyzed that account is taken of any minor maintenance work that is done, such as changing the oil in the wheel motor. For example, Figure 5 illustrates that the actual transition path of oil measurements was from A to B to C to D. If we did not account for the oil change, then the modeling process would assume that the transition was A to B to D. This would be misleading and would tend to overestimate the risk of failure.



Take in Figure 5 Oil Changes

In the EXAKT data preparation phase, the model is told what covariate values should be associated with those minor corrective events, such as an oil change (OC).

Figure 6 shows ‘missing’ or ‘irregular’ oil changes and obvious gaps. Oil ages of 7000-8000 hours are indicated which is quite unlikely with the use of mineral oils in this application. The site changed to synthetics about two years earlier to eliminate the need

for regular oil changes. However most histories, containing missing oil changes, occurred prior to 1997. It was thought that this information needed to be recovered from the commercial laboratory's files. Unfortunately these files, too, were incomplete and inconsistent with the dates and working ages in the work order database.

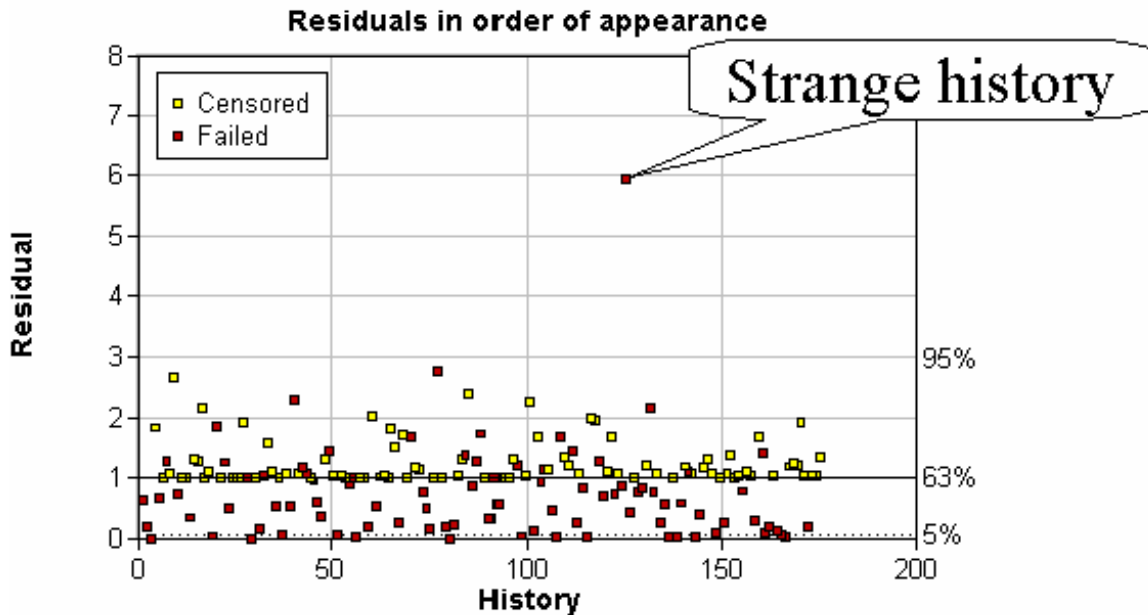
5501L 1										
Ident	Date	WorkingAge	HN	P	Event	PrevSed1	CorrSed1	PrevSed2	CorrSed	LogSed
5501R	19/10/1993	52781	1	0	*	10	470	10	470	09401642
5501R	06/11/1993	53048	1	0	*	470	950	470	950	06254539
5501R	06/11/1993	53048	1	1	OC	950	10	950	10	27279837
5501R	21/11/1993	53295	1	0	*	10	1605	10	1605	89450671
5501R	05/12/1993	53538	1	0	*	1605	2600	1605	2600	26544865
5501R	05/12/1993	53538	1	1	OC	2600	10	2600	10	27279837
5501R	27/03/1994	54028	1	0	*	10	1575	10	1575	27041782
5501R	13/04/1994	54311	1	0	*	1575	2015	1575	2015	62919126
5501R	17/04/1994	54329.75	1	1	OC	2015				
5501R	24/04/1994	54476	1	0	*	10				
5501R	01/05/1994	54609	1	0	*	1565	15433	1565	15433	14701755
5501R	17/05/1994	54839	1	0	*	15433	545	15433	545	97574491
5501R	09/06/1994	55103	1	0	*	545	190	545	190	42804663
5501R	17/06/1994	55181	1	1	OC	190	10	190	10	27279837
5501R	26/06/1994	55377	1	0	*	10	455	10	455	80951439
5501R	26/06/1994	55377	1	1	OC	455	10	455	10	27279837
5501R	06/07/1994	55529	1	0	*	10	530	10	530	02124194
5501R	13/07/1994	55661	1	0	*	530	480	530	480	27010576
5501R	06/08/1994	55931	1	0	*	480	365	480	365	33340137
5501R	21/08/1994	56153	1	0	*	365	280	365	280	66933375
5501R	06/09/1994	56461	1	0	*	280	525	280	525	21273771
5501R	26/09/1994	56723	1	0	*	525	920	525	920	03625531
5501R	12/10/1994	57005	1	0	*	920	860	920	860	50442773
5501R	03/11/1994	57309	1	0	*	860	640	860	640	45692067
5501R	21/11/1994	57583	1	0	*	640	920	640	920	03625531
5501R	21/11/1994	57583	1	1	OC	920	10	920	10	27279837
5501R	11/12/1994	57828	1	0	*	10	720	10	720	13728495
5501R	23/12/1994	58046	1	0	*	720	760	720	760	35786169

Take in Figure 6 Missing Oil Change Events

Happily, it was determined that these 'missing' oil changes did not significantly affect the model since they were relatively few in number with respect to all of the known oil changes. That is, there were a sufficient number of known oil changes in the database for the model to account for their effect on the measured data.

Building the Proportional Hazards Model (PHM)

After all the obvious data errors were eliminated or corrected, the proportional hazard model was generated. As illustrated in equation 1, the hazard function is the risk of failure of an unfailed unit at a given point in time. It is a function of both working age and the “significant” condition data. By following the iterative procedure in the EXAKT software, which is based on Cox (1972), the insignificant covariates are removed to arrive at some potential models. Those models are then tested to see how well they represent the actual data. One of the methods used by the software is known as residual graphical analysis.



Take in Figure 7 Residual Graph

Each point on the residual graph of Figure 7 represents a history, namely the time from wheel motor installation to its removal. The sample used to build the model consists of many histories. The graph shows an unusual point that is well above the 95% upper limit. This leads one to investigate the underlying data corresponding to this residual. It was discovered that some 'unusual' data were included in that history which appear to violate the model we are attempting to build.

The unusual residual value was identified as corresponding to one history from wheel-motor 5509R, with beginning event at 48900 hrs and EF (ending with failure) event at 72005 hrs.

The 'offending' data is shown on Figure 8

No events support this jump in values

Ident	Date	WorkingAge	Al	Cr	Cu	Fe	Ni	Ti	Pb	Si
5509R	2/23/94	62926	7	0	1	27	0	1	0	56
5509R	3/08/94	63171	11	0	2	43	0	0	0	63
5509R	3/27/94	63443	16	0	2	63	0	0	1	85
5509R	4/15/94	63733	26	0	2	97	0	1	1	133
5509R	5/01/94	64009	7	0	1	33	0	0	3	48
5509R	5/13/94	64244	9	0	1	52	0	0	3	59
5509R	5/30/94	64508	616	8	5	1890	17	31	8	900
5509R	6/16/94	64755	900	17	10	2500	40	47	12	900
5509R	7/05/94	65023	573	12	6	2500	29	28	8	900
5509R	7/19/94	65268	572	11	6	2500	31	26	7	900
5509R	8/07/94	65497	226	6	4	1509	16	12	5	900
5509R	8/21/94	65533	236	4	4	1003	8	12	4	879
5509R	9/08/94	65804	56	1	1	415	3	3	1	319
5509R	9/23/94	66252	83	3	1	581	6	6	2	444
5509R	10/12/94	66561	31	2	1	446	2	8	1	309
5509R	10/23/94	66750	32	2	1	554	4	8	1	343
5509R	11/12/94	66933	17	0	1	246	1	3	0	164
5509R	11/30/94	67334	26	2	1	430	5	5	1	265
5509R	12/23/94	67587	24	1	2	321	3	3	2	216
5509R	1/11/95	67886	23	0	2	111	0	1	2	111
5509R	1/26/95	68102	6	0	0	121	0	1	0	71
5509R	2/09/95	68350	11	0	1	183	0	1	1	101
5509R	2/27/95	68624	7	0	1	103	1	1	1	60
5509R	3/12/95	68858	8	0	0	144	0	1	1	80
5509R	4/29/95	69174	13	1	0	260	2	1	3	114
5509R	5/10/95	69365	15	3	1	345	6	1	4	127
5509R	5/29/95	69666	7	1	1	127	3	2	2	60
5509R	6/13/95	69944	10	0	1	236	1	3	1	106
5509R	7/12/95	70415	6	0	1	177	1	1	1	69
5509R	8/14/95	70963	4	0	1	108	1	1	1	47
5509R	9/14/95	71457	2	0	0	72	0	0	0	38

Take in Figure 8 Investigating the strangeness

The Fe values in the left-circled region of Figure 8 have an inexplicable pattern. Fe jumps to high values, but truncated at 2500 PPM due to instrument saturation, and remains in the same range for a few more inspections. Then, the readings fall back to low values. No events were recorded to explain these sudden jumps.

Having no event data to support such high values of Fe and Si, the model was regenerated and the fit tested again after removing that history from the working data set. Statistical and graphical goodness-of-fit testing procedures are applied by the software as part of the modeling procedure. The model's fit to the data improved immediately. The model building algorithms no longer had to accommodate obviously contradictory and misleading information.

The forgoing describes data problems which were encountered and which were relatively easy to correct using the statistical and graphical tools available in the software's functionality.

However a different (and more fundamental) problem occurred regarding the definition of wheel motor failure. These units seldom failed "functionally". Being monitored monthly, there were few "in-service" failures requiring that a haul truck be removed immediately from operation. Nevertheless the model required an objective determination

that a unit had failed. To provide the event data to the model, it was necessary to go through the past work order records to find the failures and the preventive removals. Initially, the tradesman remarks were used for this purpose, such as "High iron in oil sample and high hours, removed and replaced wheel motor." This event was then classified as a "failure". However, on reviewing the re-builder's report attached to each invoice it became clear that some events initially classified as a failure should be treated as a suspension. For example: If the gears had been replaced because they failed an ultrasonic test or they were obviously in a failed state then that would be classed as a failure. But if they were replaced simply because it was expedient to do so, or if the unit was only generally rebuilt with no real internal damage, then that was considered a suspension.

With the definition of suspension and of failure thus clarified, a proportional-hazards model was found which was shown to be a "good fit".

The model containing the covariates iron and sediments was found to be good, both by graphical residual analysis and by the Kolmogorov-Smirnov statistical test applied automatically by the software. The results of the analysis are displayed in Figure 9. Covariate significance is tested by the Wald statistic, the square of the standardized estimate of the parameter which follows a chi square distribution with 1 degree of freedom. (Note: A few missing sediment values had been replaced by the values from previous inspections prior to the analysis, hence the reason for using the notation CorrSed in Fig. 9). The PHM is thus:

$$h(t) = \frac{2.891}{23360} \left(\frac{t}{23360} \right)^{1.891} e^{0.002742*Fe+0.0000539*Sed}$$

PHM Parameter Estimation
PHM_CRC (FeCorrSed(Opt3)) (Wheel Motors)

Summary of Events and Censored Values

Sample Size	Failed	Censored (Def)	Censored (Temp)	% Censored
175	88	41	47	49.7

Summary of Estimated Parameters (based on ML method)


Parameter	Estimate	Sign.(*)	Standard Error	Wald	DF	p - Value	Exp of Estimate	95 % CI	
								Lower	Upper
Scale	2.336e+004		1057					2.129e+004	2.544e+004
Shape	2.891	Y	0.2692	49.33	1	0		2.363	3.419
Fe	0.002742	Y	0.0002274	145.4	1	0	1.003	0.002296	0.003188
CorrSed	5.395e-005	Y	2.476e-005	4.747	1	0.02934	1	5.42e-006	0.0001025

(*) Based on 5% significance level. Shape = 1 tested, Gamma (Cov) = 0 tested.

Take in Figure 9 The Proportional Hazards Model

Obtaining the Decision Model


After determining the PHM we are ready to establish the optimal decision model (see Jardine et al 1997) that incorporates economic considerations along with the risk estimate obtained from the PHM. The decision policy was calculated with a cost ratio of 3:1 (\$20K for preventive replacement cost, \$60K for failure replacement cost, based on the invoices of past repairs by outside contractors). See Figure 10. The cost ratio will vary amongst applications. It could include the costs associated with operational consequences depending on current production and coal market conditions. The model includes the effects of regular maintenance intervals (oil changes) at 500 hrs that occurred regularly during most of the histories prior to the changeover to synthetic oil.



$$Z = 0.002742 * Fe + 5.3955e-005 * CorrSed$$

A model calculated without including the oil change intervals would tend to underestimate predicted failure times (Note: When data was used for estimation, covariate values were set to zero on oil change, as defined in the CovariatesOnEvent table. The PHM parameters and transition probabilities are then estimated from this adjusted data. So, in the subsequent calculation of averages and costs, we have to take into account that covariates will regularly have their values reduced. If not, covariates will reach high values in the calculation process faster than in real data, and thus produce a higher estimated risk of failure than is really the case).

At present the model does not attempt to optimize inspection frequency (a future research feature), leaving the decision up to the user. It is to be noted that no operational savings were accounted for. This was due to the present unfavorable coal market conditions causing the mine to operate below its capacity. It is expected that as market conditions improve higher cost ratios would be used. Current strip ratios (total material removed versus sellable material) would also affect the savings associated with increased availability and reliability of the units. Figure 11 demonstrates the sensitivity of the overall savings to changes or inaccuracies in the cost ratio.



Replacement Cost Ratio (C+K)/C	Optimal Hazard Level	Ratio of Optimal Costs [%]	Increase Over Optimal Cost Using **, [%]
1	No Limit	44.76	147.12
2	0.000122139	80.48	103.04
* 3 *	7.16861e-005	100.00	100.00
4	5.39354e-005	113.94	102.75
5	4.26086e-005	126.59	105.97
6	3.87334e-005	135.05	111.98
7	3.50239e-005	143.42	117.35
8	3.13143e-005	151.79	122.13
9	2.76048e-005	160.15	126.41

Take in Figure 11 Sensitivity Testing

In real situations, the actual ratio of failure and preventive replacement costs may not be well known. Furthermore the dynamics of industry are such that costs can change with changing technology, production, and market conditions. Therefore one would like to know, to what degree the true total cost per unit time and the optimal policy would change with changes in cost ratio. The software enables sensitivity analysis to be undertaken and generates a graph and corresponding tabular data, see Figure 11.

The curves on the graph are interpreted as follows. Solid Line: If the actual cost ratio (CR) differs today from that specified when the model was built, that means that the current policy (as dictated by the Optimal Replacement Graph of Figure 10) may no

longer be optimal. The line indicates the increase in percentages that will be incurred above the optimal cost/unit time for the actual cost ratio by adhering to the current policy. For example, if the actual cost ratio is 5 and we are using a model which is based on CR=3, then the increase in the cost incurred by following that (wrong) policy is around 6% (5.98). In other words the solid line represents the sensitivity of costs to changes in CR. Dashed Line: Again, assume the actual cost ratio has strayed from what was used when the model was built. If the model is rebuilt using the new ratio the dashed line tells how much the new optimal cost would differ from that of the original model. (Note that the sensitivity graphs assume that only C_f (failure replacement cost) changes and C_p (preventive replacement cost) remains unchanged.) In other words the dashed line represents the sensitivity of the optimal policy to changes in CR. The graph indicates that moderate overestimation of the cost ratio does not significantly affect the average long run cost but provides a more conservative policy from the point of view of risk of failure.

CR=3	Cost [\$/hr]	Preventive Repl. Cost [\$/hr]	Failure Repl. Cost [\$/hr]	Prev. Repl. [%]	Failure Repl. [%]	Expected Time Between Replacements
Optimal Policy	2.26497	1.1047 (48.8 %)	1.16027 (51.2 %)	74.1	25.9	13409.7
Replacement Only At Failure	3.04139	0 (0.0 %)	3.04139 (100 %)	0.0 %	100.0 %	19727.8
Saving	0.776415 (25.5 %)	-1.1047	1.88111	-74.1 %	74.1 %	-6318.13
CR=5	Cost [\$/hr]	Preventive Repl. Cost [\$/hr]	Failure Repl. Cost [\$/hr]	Prev. Repl. [%]	Failure Repl. [%]	Expected Time Between Replacements
Optimal Policy	2.83999	1.60461 (56.5 %)	1.23537 (43.5 %)	86.7	13.3	10800.9
Replacement Only At Failure	5.06898	0 (0.0 %)	5.06898 (100 %)	0.0 %	100.0 %	19727.8
Saving	2.22899 (44.0 %)	-1.60461	3.83361	-86.7 %	86.7 %	-8926.89
CR=6	Cost [\$/hr]	Preventive Repl. Cost [\$/hr]	Failure Repl. Cost [\$/hr]	Prev. Repl. [%]	Failure Repl. [%]	Expected Time Between Replacements
Optimal Policy	3.0595	1.9303 (63.1 %)	1.12921 (36.9 %)	91.1	8.9	9440.65
Replacement Only At Failure	6.08277	0 (0.0 %)	6.08277 (100 %)	0.0 %	100.0 %	19727.8
Saving	3.02327 (49.7 %)	-1.9303	4.95357	-91.1 %	91.1 %	-10287.2

Take in Figure 12 Potential Savings


The cost analysis summary shown on Figure12 indicates a saving of 25%, when CR=3, over the “replace only on failure” (ROOF) policy, whose costs approximate those of the site’s past policy.

Decision model results are also calculated for cost ratios of 5 and 6. As the cost ratio increases we can observe an increase in both the optimal policy cost as well as an increase in savings. The optimal decision models in these cases indicate more frequent

preventive replacements (from 74% to 91%) will result from applying the optimal decision policy in order to avoid costly failures. (Note: There is a slight discrepancy between the expected time between replacements for the ROOF policy, when CR=3 and CR=5 and 6. This is due to the numerical calculation procedure.)

APPLICATION OF DECISION MODEL

Once the decision model was built (see Figure 10), data was examined from previous histories to see what the decision model would have recommended for situations in which the wheel motor failed. One illustration of such a history is shown in Figure 13. This graph provides a recommended decision based on inspection data (covariates and working age).



Take in Figure 13 The Decision Graph

The decision ‘Replace immediately’ was suggested by the model for the first time at the inspection point at working age = 11384 hrs, 286 hours (about two weeks) prior to failure (reported at 11660 hrs). The following inspection at working age = 11653 hours, 7 hours prior to failure, also suggests the replacement of the wheel motor. The first warning may have been sufficient, given sample turnaround time of 48 hours, to prevent the consequences of failure. Even prior to 11384 hours it can be seen from the decision graph that the results of the measurements indicate that a replacement recommendation was imminent. Note that the zero points on the graph indicate default measurement values of zero (imputed by the software) immediately following oil changes.

The economic benefit associated with basing the maintenance policy on the Decision Policy Graph (Figure 10) is identified through this type of investigation. This analysis indicates a potential saving of between 20%-30% compared to the current practice.

CONCLUSION

The case study first demonstrates the value of using the technique of PHM to assist maintenance professionals to interpret condition data by identifying the key risk factors and their relative influence on the health of equipment in general, and wheel motors in particular. Economic considerations were then blended with the PHM risk model to identify the optimal decision chart. The study then indicates that the implementation of a condition-based maintenance strategy based on the decision chart would result in a cost reduction of 20%-30%.

REFERENCES

1. Cox, D.R., (1972) "*Regression models and life tables (with discussion)*", J.Roy. Stat. Soc. B, Vol. 34,pp. 187-220..
2. Jardine, A.K.S., Banjevic D. and Makis V, (1997) "*Optimal replacement policy and the structure of software for condition-based maintenance*", Journal of Quality in Maintenance Engineering, Vol. 3, No.2, pp. 109-119.
3. Campbell, J.D. and Jardine A.K.S. (Editors), (2001) *Maintenance Excellence: Optimizing Equipment Life-Cycle Decisions*, Marcel Dekker, (Chapter 12: Optimizing Condition Based Maintenance, by M. Wiseman).