CENTRE FOR MAINTENANCE OPTIMIZATION AND RELIABILITY ENGINEERING

Update on Interactions with Consortium Members and Industry Partners

June 6, 2017

DEPARTMENT OF MECHANICAL AND INDUSTRIAL ENGINEERING

UNIVERSITY OF TORONTO

5 KING'S COLLEGE ROAD

TORONTO, ONTARIO, CANADA M5S 3G8

T: +1 (416) 978-2921 F: +1 (416) 946-5462

cmore.mie.utoronto.ca

Table of Contents

C-MORE News

Neil Montgomery, C-MORE

C-MORE Director

In April, 2017, Mike Kim informed the department that he has accepted an appointment at the University of British Columbia beginning July 1 and will no longer serve as C-MORE Director.

The department has begun the process of attempting to a replacement.

The NSERC CRD grant with TTC and Toronto Hydro will be transferred to another professor in the department until a new Director is appointed.

The remainder of this document will focus on our work with companies since the last full Consortium meeting.

C-MORE activities with collaborating companies

Since December 2016, C-MORE lab members have been working on research and meeting with collaborating companies. C-MORE is currently involved in the following projects with industry partners:

- Teck: Neil Montgomery and Dragan Banvevic are working with David Williams and Scott Hansen on a problem involving estimating reliability of series systems in the absence of any component failres.
- TTC: Neil Montgomery completed work on relays maintenance and inspection interval study and a track switch reliability and maintenance case study. He is near the completion of a train stop reliability case study.
- Toronto Hydro: We plan to expand and generalize prior work done on asset hazard function estimation problems.
- Canadian Bearings: C-MORE and Canadian Bearings have completed prototype work on a project to enable Canadian Bearings to offer a simplified criticality analysis to its customers.

In addition, Neil Montgomery and Andrew Jardine visited DICTUC Tribologia in April, 2017, to deliver training on the use of EXAKT focusing on oil analysis applications in the Chilean mining sector. During this training session we discovered a bug relating to the computer's internal date format. Programmer Dustin Wehr fixed the problem and we can make available an updated bug-fix version of EXAKT for Consortium members.

Technical Reports

Teck Mine Hoist Reliability Problem

Neil Montgomery and Dragan Banjevic

Problem description

We received the following question from David Williams and Hansen from Teck.

"I was wondering if I could get some help. We are trying to justify modernizing the drive and PLC for our mine hoist. The system was installed in 2005 and we are finding that many key components are no longer supported. None of the key components have failed but as I've had the team go through and identify critical spares (something that hadn't been done until recently), we're finding that the only source for many of the parts is eBay. This poses a huge risk as there are no guarantees to the condition of the components, or that they will be available the next time we look.

Countering this is a current mine life of 3 years. I've attached a spreadsheet that I had my process control engineer put together in an effort to quantify the probability of a catastrophic PLC or drive failure. Could you please take a look at our methodology and let me know if you think there would be a better way to evaluate the problem, and if there would be a better way of incorporating declining parts availability over time."

Details

The spreadsheet we were given contained three calculations in which MTBF of 12, 24, and 60 years were given to each component in two sub-systems of the hoist. The components were mostly electronic in nature (control systems, power supplies, etc.) The hoist has other components, but these are the ones that were difficult to obtain.

All parts were listed as "repairable". Most were difficult to obtain. Teck was concerned with the reliability of parts obtained via eBay, so their failure rates were adjust upward by 10% to account for their possible unreliability.

The "Drive" subsystem had 8 largely identical parts. The "PLC" subsystem had 21 components which were not all identical.

Teck converted component MTBFs to failure rates and added them to obtain a system failure rate, which they then used to compute reliability and unreliability for 1, 2, and 3 year planning horizons.

The calculations themselves were performed correctly. However, we has some comments about their underlying assumptions.

Our Commentary

The calculations done in the spreadsheet are correct, as far as their intentions go. Specifically, if all the components have either 12, 24, or 60 year MTBFs, then all the subsequent calculations are correct.

There are a few underlying assumptions behind the calculations. Some are probably reasonable, such as the constant failure rate assumption being used for each system and all their components. Others might not be, such as the assumption that all components have equal rates. I think the impact of the derate factor probably ends up being very small.

On the other hand, I would consider the implication of giving 12, 24, and 60 year MTBF to each and every component in the two systems. An implication of this is included in the spreadsheet you sent, but it might have not been obvious because the interim calculations were all done in hours. If you look on the 12 year tab, the Drive Whole System MTBF is given as 12888 hours, which is only 1.47 years. The probability of the system having lasted 12 years already with no failures (which it has), and with such a low MTBF (implied by giving the 8 components each an MTBF of 12 years) is really low (0.000287). The situation is even more extreme when you look at the PLC, with so many more components, resulting in a PLC Whole System MTBF of 0.51 years. There's just no way a system with that MTBF would have lasted 12 years the way it has with no failures.

I understand that 12 was used as a pessimistic assumption - but it might not have been clear just how pessimistic it was. Even 24 and 60 years, on a per component basis, are pessimistic (severely so in the PLC case.) I added a tab to the spreadsheet "System-Based". The first table shows the implications on system reliability of giving the same 12 / 24 / 60 year MTBF to each component.

Another way to think of the problem could be as follows. The system survived 12 years. Was this surprising? Thinking back to 12 years ago, what was the chance we would have lasted this long with no failures among these components? Perhaps there was only a 10% chance - in other words, we're really lucky nothing failed. Or maybe it was 50-50. Or maybe there was a 90% chance that none of these components would have failed.

Implications of assigning same MTBF to each component

	System MTBF in Years			**Chance of having lasted 12 years**		
Per-Component MTBF in years	12	24	60	12	24	60
Drive	1.47	2.94	7.36	0.0287%	1.6938%	19.5678%
Hoist	0.51	1.02	2.55	0.0000%	0.0008%	0.9002%

The second table in the new tab translates this range of feelings (we were lucky at 10%, all the way up to 90% chance of having lasted this long) into the system-wide MTBF, system-wide rate, and then 1, 2, and 3 year projections of system-wide survival probabilities. If there was about a 50-50 chance of making it this far with no failures, that implies an MTBF of 17.31 years, for example, and the probabilities of lasting a further 1, 2, and 3 years, are 94.4, 89.1, and 84.1 percent, respectively. Note that even if you were very lucky, with only a 10% chance of having made it so far with no failures. You still end up with a 56% chance of lasting the next three years.

The system-wide "How fortunate have we been" approach					
	Chance that we would have lasted these 12 years?				
	10%	30%	50%	70%	90%
Implied System MTBF	5.21	9.97	17.31	33.64	113.89
Implied System Rate	0.1919	0.1003	0.0578	0.0297	0.0088
Chance of surviving one more year	82.5%	90.5%	94.4%	97.1%	99.1%
Chance of surviving two more years	68.1%	81.8%	89.1%	94.2%	98.3%
Chance of surviving three more years	56.2%	74.0%	84.1%	91.5%	97.4%

It will be very hard in any kind analysis to get MTBFs and rates for the individual components themselves. The two extremes are:

- that the rates are all equal, in which case you divide the system rate by the number of components, resulting in very long MTBFs for each part.
- that one component completely dominates, in the sense that its rate is essentially the system rate and all the others are essentially 0.

No component could have a failure rate higher than the implied system-wide rate.

It's hard to say what actual part stocking strategy you might undertake from the analysis. If the system failure probability projections I've made are acceptable risks to take, you could do nothing, especially considering that it seems the parts are at least repairable (unless I misunderstand the "Repair from 3rd Party" column). If the risks are unacceptable, it might be a matter of hedging your bets by prioritizing some combination of:

- those parts which are more likely to actually fail (if known)
- those parts more difficult to procure (if known)
- those parts more difficult to repair (if known)

with the understanding that your could hedge as best you can and still lose the bet. As I mentioned, the derate consideration has almost no impact on overall reliability, if we're only talking a 10% change to a component's failure rate.

TTC Relays Case Study

Neil Montgomery, C-MORE

Executive summary

The TTC's has a preventive overhaul program for its relays on either a 5 year or a 10 year interval. The Centre for Maintenance Optimization and Reliability Engineering assessed the program using available data from work order and fault records and have determined the following:

- 86% of all overhauls start within 30 days of their overhaul due dates, which would be considered indicitave of good performance of this aspect of the maintenance program.
- since most overhauls are performed close to their scheduled times we would not be able to determine if a different interval would lead to better results.
- we found that relays that had ever experienced an overhaul delay (in excess of 100 and 200 days) were overrepresented in the population of relays that had ever failed.
- the manufacturer and possibly the age of the relay is also related to the likelihood of failure and may explain some of the phenomenon from the previous point.
- nevertheless, the TTC should continue to carefully control the preventive maintenance program to ensure no significant delays to relay overhauls.
- the manufacture dates of many of the relays are only known up to a possible range of dates, in some cases quite wide, making the application of many traditional maintence models difficult. We will make a recommendation about data collection practices to address this issue.
- the location history of the relays is not kept, so it is not possible to perform a data-driven analysis based on differences in workload by location.
- for each manufacturer, a raw failure rate seems to increase with the number of overhauls performed. It remains an item of discussion as to what action, if any, to take based on this information.

Problem overview

A relay is an electromechanical device that implements the safety logic conditions for the trains. TTC relays are overhauled on either a 5 year or 10 year time interval and otherwise inspected and repaired (if necessary) whenever a relay fault is suspected.

The purpose of this study is to evaluate the evidence of the overall effectiveness of the 5 and 10 year intervals, as these overhaul policies consume a substantial percentage of the relay shop's resources, and the policies impact the unobstructed operation of the subway service.

The data files

We received four datasets relating to relays.

- Master list: `Relay_Assets_v2.xlsx` consisting of relays currently in operation.
- Preventive maintenance details `Relay_PM_Details_v2.xlsx` consisting of the PM tasks defined associated with each relay.
- Work orders: `Relay_PM_WO_1997-2016_v2.xlsx` consisting of PM work records for relays. There are 31982 such records.
- Faults: `Relay_Fault_RS_1997-2016_v3.xlsx` consisting of work orders relating to suspected relay faults. There are 974 relay faults recorded.

We performed a preliminary analyis of these files to identify obvious problems, such as missing or unaccounted for relays. A small number of anomalies were found and TTC staff answered our questions to our satisfaction, with most relays identified as either scrapped or non-vital (which are not overhauled) and not important to the analysis.

It is important to note that we have not undertaken a detailed record-by-record anaysis of the free-form text fields in the various files - a task which could take many months and is not usually necessary for population-based analyses such as done in this report. There are probably work order records indicating the relay was in a failed state and there are probably a few records in the fault files when the relay was healthy. A basic overall assessment indicated that such errors do not affect the overall analysis.

If we were to undertake a detailed relay-by-relay asset-based analysis then each record would indeed need to be assessed. We will further investigate the relays records as a first step towards determining if a deeper analysis is warranted.

Relay ages

None of the files contained the manufacture date of the relays. The best available information about the relay ages is expert knowledge based on relay manufacturer. We have added manufacture dates to the data from which approximate ages can be inferred according to the following:

Manufacturer	Year Range	Date Used
GRS 927 Series and AC Vane	2003-2004	2003-01-01
Other GRS	1963-1968	1965-01-01

SGE	1954	1954-01-01
US&S	1973-1974	1973-01-01
WABCO	1973-1974	1973-01-01
Transcontrol	1978-1996	1985-01-01
Safetran	1996	1996-01-01
Alstom	2002	2002-01-01

Note in particular the wide range of possible ages for the Transcontrol relays, which really cannot then be used for any subsequent analysis for that manufacturer.

Relay locations

TTC and C-MORE discussed the possibility of also examining the relay locations, since some are considered more stressful than others. However, the vast majority of all locations (97% in the work orders dataset and 90% in the faults dataset) list the relay location as being in the relay shop, so no data-driven analysis is possible on this question. It is possible that historical relay locations are available in the databases but we did not query the data to extract them. We will further investigate the availability of relay location data.

Timeliness of relay overhauls

The main focus of the study concerns the effectiveness of the 5 and 10 year overhaul program for the relays. From C-MORE's point of view, we will look at how many overhauls are actualy performed on time and the relationship between timeliness and reliability.

The following table shows the percentages of overhauls commenced within 30 days of their scheduled dates, broken down by 5 year and 10 year overhaul schedules.

Interval	Count	Within 30 Days
1825 DAYS	11890	81.1%
3650 DAYS	15781	90.4%

The percentage of preventative maintenance actions completed on time is a maintenance performance mangement metric. In our experience these observed percentages are indicitive of good performance. (We are not able to benchmark against other subway systems specifically.)

Note on "optimizing the interval"

Industry often seeks to use data to determine what the best preventive maintenance (PM) interval *should* be, given the data. The answer depends on the organization having enough variation in actual PM interval times to determine a relationship between PM interval and failure rate. In the case of the TTC because of their timely overhaul completion for relays, there is not enough variation to draw a conclusion about whether a longer or shorter interval would be better.

Delayed overhauls and failed relays

Not all overhauls are commenced on schedule. As noted earlier, there have been a total of 974 relay faults recorded in the data. (It is also important to note that there were surely many faults and certainly many overhauls over the years the predate the current databases.)

Is there a relationship between delay and relay reliability? We looked at those relays that had *ever* suffered overhaul delays in excess of 50, 100, and 200 days at any time in its recorded history. If these delays didn't matter, we would expect them to appear in the faults records in the same percentage as those which did not. But we observed the following:

Ever delayed by?	percentage of all relays	percentage among failed relays
50 days	16%	29.4%
100 days	8.7%	18.4%
200 days	4.6%	10.4%

An incidence of significant delay in a relay's overhaul history is correlated with the incidence of failure. The data do not show a clear and immediate causation. The failures that occurred subsequent to a delayed overhaul sometimes occurred well after the delay.

It is possible that relays that are older would therefore have greater chances of ever suffering a delay and also ever failing, providing a simple explanation for this phenomenon.

In our case, the age of the relay is primarily related to its manufacturer. It is not possible to determine the difference in effect of age versus manufacturer. Certainly in the case of the Transcontrol relays, the ages are not sufficiently known to draw any conclusion. The possible age range for the other relays has probably been underestimated.

The following complex table summarizes the situation.

Manu.	Mean failed age	Mean nonfailed age	perc. of all	perc. of failures	perc. of delayed 50	perc. of delayed 100	perc. of delayed 200
ALSTOM	11.016	11.537	0.2%	0.3%	0.0%	0.0%	0.0%
GRS_New	9.282	9.110	27.5%	6.4%	17.7%	17.9%	19.6%
GRS_Old	48.363	47.810	17.0%	43.5%	34.1%	35.1%	36.2%
SAFETRAN	17.087	16.446	5.5%	2.0%	5.0%	4.1%	4.3%
SGE	59.640	58.798	6.2%	11.5%	13.8%	15.9%	16.8%
TRANSCO	27.067	25.360	16.9%	18.3%	12.2%	10.2%	8.0%
US&S	38.969	37.909	12.2%	3.2%	2.6%	1.9%	1.3%
WABCO	38.958	37.258	13.5%	13.0%	13.8%	14.2%	13.0%

Some explanation is required. The GRS relays have been split into the new ones of 2002/3 vintage and the originals. The mean failed and non-failed ages are there to give an indication of the ages of the relays for each manufacturer (with the TRANSCO ages not to be taken too seriously). Each percentage column has the percentage of relays satisfying the indicated condition. So "perc. of all" means percentage of all relays; "perc. of faulures" means percentage of failures, etc., so that the percentage *columns* add to 1.

We can see that the manufacturers are not represented among the failures in percentage to their numbers in the population, but it is not simply a function of relay age.

(Note that TTC and C-MORE did not specifically discuss the overall performance characteristics of the relay manufacturers. Some of these numbers may not come as a surprise to TTC.)

We conclude therefore that the overrepresentation of relays that had suffered at some point a delay in being overhauled can be partially, but not entirely, explained by the age/manufacturer (in this dataset these two characteristics are difficult to separate). The TTC should continue to ensure that overhauls are commenced without significant delay.

The following table is similar to the previous, except that the percentages are within each manufacturer row.

Manu.	Mean failed age	Mean nonfailed age	Perc. of all	Perc. fault	Perc. delayed 50	Perc. delayed 100	Perc. delayed 200
ALSTOM	11.016	11.537	0.2%	12.0%	0.0%	0.0%	0.0%
GRS_New	9.282	9.110	27.8%	1.3%	9.4%	5.3%	3.2%

GRS_Old	48.363	47.810	17.2%	14.6%	29.1%	16.8%	9.5%
SAFETRAN	17.087	16.446	5.5%	2.1%	13.4%	6.0%	3.6%
SGE	59.640	58.798	6.3%	10.5%	32.3%	20.8%	12.1%
TRANSCO	27.067	25.360	17.1%	6.2%	10.5%	4.9%	2.1%
US&S	38.969	37.909	12.3%	1.5%	3.1%	1.3%	0.5%
WABCO	38.958	37.258	13.7%	5.5%	14.8%	8.6%	4.3%

Number of overhauls and reliability

We would also like to examine the relationship between the number of overhauls and relay reliability. There are several ways to approach this problem. In our case we do not have very detailed information about relay age in many cases. We will take a simpler approach, which will simply be to count the number of overhauls in the data for each relay, and count the number of failures occuring during after each number of overhauls.

There are two complications in the data. The first we have mentioned already. We have no information about overhauls and failures before the beginning of the datasets, and surely there are events that are missing.

The second is that there are many entries in the work order datasets that indicate a scheduled overhaul far sooner than should have been indicated. This point should be discussed and analyzed further. For now we have assumed that two work order records occurring within one calendar year of other refer to the same general overhaul event with multiple work orders relating to different tasks within the same overhaul.

We break down the results by relay manufacturer, cutting off the table when the sample size becomes too small. There are too few Alstom relays to report on at all.

GRS "new" relays:

Number of overhauls	Non-failed	Failed	percentage
0	7941	14	0.18%
1	4296	32	0.74%
2	360	9	2.44%

GRS "original" relays:

Number of overhauls	Non-failed	Failed	percentage
0	4814	45	0.9%
1	2734	36	1.3%

2	2142	119	5.3%
3	1436	134	8.5%
4	842	75	8.2%
5	183	24	11.6%

SAFETRAN:

Number of overhauls	Non-failed	Failed	percentage
0	0	3	100%
1	860	11	1%
2	496	5	1%
3	80	0	0%

TRANSCO:

Number of overhauls	Non-failed	Failed	percentage
0	0	6	100.0%
1	2656	57	2.1%
2	1998	71	3.4%
3	201	35	14.8%

USS:

Number of overhauls	Non-failed	Failed	percentage
0	0	5	100%
1	1899	13	1%
2	296	10	3%

WABCO:

Number of overhauls	Non-failed	Failed	percentage
0	0	10	100%
1	2117	18	1%
2	1789	59	3%
3	506	23	4%
4	198	14	7%

We also report for the SGE relays, although there is no choice but to refurbish these relays since replacements are not available:

Number of overhauls	Non-failed	Failed	percentage

0	0	2	100%
1	983	21	2%
2	871	42	5%
3	363	34	9%
4	177	15	8%

In each case the rate of failure increases with the number of overhauls. The failure rates themselves remain low. We have not yet attempted to analyze the tradeoffs involved in using a different asset management strategy for relays using this failure rate information. It remains an open question how to best use this information.

Recommendations

- A global "begin" date should be recorded for all assets. This could be a manufacture data or an original installation date. This recommendation is not specific to relays. It should be a policy for all assets. The age of an asset needs to be known in order to use almost any reliability and maintenance model.
- Precise location data on the relays should be recorded any time a relay is moved. The data should include the asset identifier, the date of the movement, and the new location.
- TTC should continue to ensure that overhauls are commenced without significant delay.

Appendix - Relay Movement Data

We received supplemental data regarding the movement of relays to and from operational locations, the repair shop, and storage facilities. Combining such movement data with reliability and maintenance data leads to some complex issues of data cleaning, which is typical across all industries in our experience. Our basic approach to the problem was to assume many relay movements over a short period of time tended to refer to one main "event", either a fault or an overhaul. The location of the fault would be assigned to the last known operational location.

There are 740 locations with faults assigned to them. Here is the distribution of locations and numbers of faults:

n_faults	Count
5	2
4	9
3	18
2	97

1	614

TTC staff can examine those locations with multiple faults to determine if there is any possible cause.

TTC Switches Case Study

Neil Montgomery, C-MORE

Executive Summary

The Centre for Maintenance Optimization and Reliability Engineering has analyzed the available reliability and maintenance data for the TTC's fleet of switch machines, including work order data, preventive maintenance data, a small sample of switch movement data, and a limited number of switch machine installation dates. We have determined the following:

- the overall reliability of the entire fleet of switch machines has steadily improved since 2007, after a decline in reliability from 1998 to 2007.
- switch machines are identified by the location, rather than by an asset number which could be used to track individual machines. We recommend that switch machines have their own asset numbers, which would allow for a much wider variety of analyses.
- the ages of most switch machines is not currently tracked, so it is not possible to assess the impact of age on failure risk.
- further to the previous two points, any asset class important enough to TTC operations should be considered for individual tracking, including installation date information.
- switches located outdoors have a small increased risk of failure, even after accounting for "external" failure causes such as ice, snow, obstruction, and debris. Otherwise we found no known risk factors for increased risk of failure.
- there are substantial differences in the number of faults/switch when grouped by manufacturer, but this may be confounded by an unknown age factor, for example.
- otherwise we found no risk factors associated with increased risk of failure.

Problem overview

The TTC operates a fleet of switch machines in its subway lines and subway train yards. The switch machines facilitate routine car movements between tracks within yards for positioning purposes, routine train movements at the ends of the subway lines to allow trains to turn around, and sporadically elsewhere on the subway lines to allow for service adjustments and other required movements. A faulty switch machine can cause a significant disruption to service.

The purpose of this analysis is to use the available reliability and maintenance data relating to switch machine locations to inform possible changes to switch machine asset management, including to identify risk factors relating to switch machine faults.

The data files

It is important to note that switch machines do not appear to be tracked (in Maximo or elsewhere) as switches, per se, using the `ASSETNUM` identifier. They are tracked using a combination of location in the `ASSETNUM` field and the entry `Switch` in the `DESCRIPTION` field. This method of tracking has a major impact in the kind of analysis that can be done.

We recommend that switch machines at least receive their own asset number for tracking, with the possibility of tracking individual components depending on the types of reliability analyses the TTC might wish to perform on its tracking switches.

For most of the rest of this report, when we write "switch" we really mean "switch machine location"

We received four datasets relating to switches.

- Master list: `SwitchMachine_Assets_rev21070421.xlsx` consisting mainly of location and manufacturer. The list identifies 368 switches. This file was updated in April, 2017, to include a new column with a specific mainline/yard indicator and a model column.
- Preventive maintenance details `SwitchMachine_PM_Details.xlsx` consisting of the PM tasks defined associated with each switch.
- Work orders: `SwitchMachine_PM_WO_1997-2016_v2.xlsx` consisting of 17667 PM work records for switches.
- Faults: `SwitchMachine_Fault_WO_1997-2016.xlsx` consisting of 3327 work orders relating to suspected switch faults, detected by transit control during operations and testing.
- Corrective maintenance actions: `SwitchMachine_CM_WO_1997-2016.xlsx` consisting of a small number (179) of corrective actions for deficiencies noted on inspection or for other reasons.
- Switch movement data: `Switch movements - average Sep 30 to Oct 6 2016 (2).xlsx` consisting of a one week sample of numbers of switch machineoperations, along with a variable stating whether the switch is outdoors or indoors. Yard switch movements are not available, and they are all outdoor switches.

- Some age data: `SwitchMachine_Assets_Install.Dates.xlsx` consisting of installation and refurbishment dates for some of the switches. For most switches there is no information

We performed a preliminary analyis of these files to identify obvious problems. No serious problems were found.

Once again it is important to note that we have not undertaken a detailed record-by-record anaysis of the free-form text fields in the various files - a task which could take many months and is not usually necessary for population-based analyses such as done in this report. In this case since switches are not individually identified there is probably no value in going to such detail.

We did identify some comments which suggest external causes for switch faults by searching for the words `ice|snow|obstruction|debris`.

Analysis of faults data

Impact of non-yard switch movements on faults

It is possible that switches could deteriorate due to a greater workload, so we examined the relationship between the number of movements and the number of faults.

It is important to note that we were working with a one week sample of movement data from which we extracted an average. This average is best understood as a ranking mechanism. It could be possible to use a larger sample of movement data, but the point is mainly to identify the switches that are often used. More movement data is likely to merely reinforce this identification rather than adding more information.

Movement data are not available for yard switches.

Several analyses were performed, but a series of simplest plots are the most revealing in this case. First is a basic plot of movement average versus number of faults.

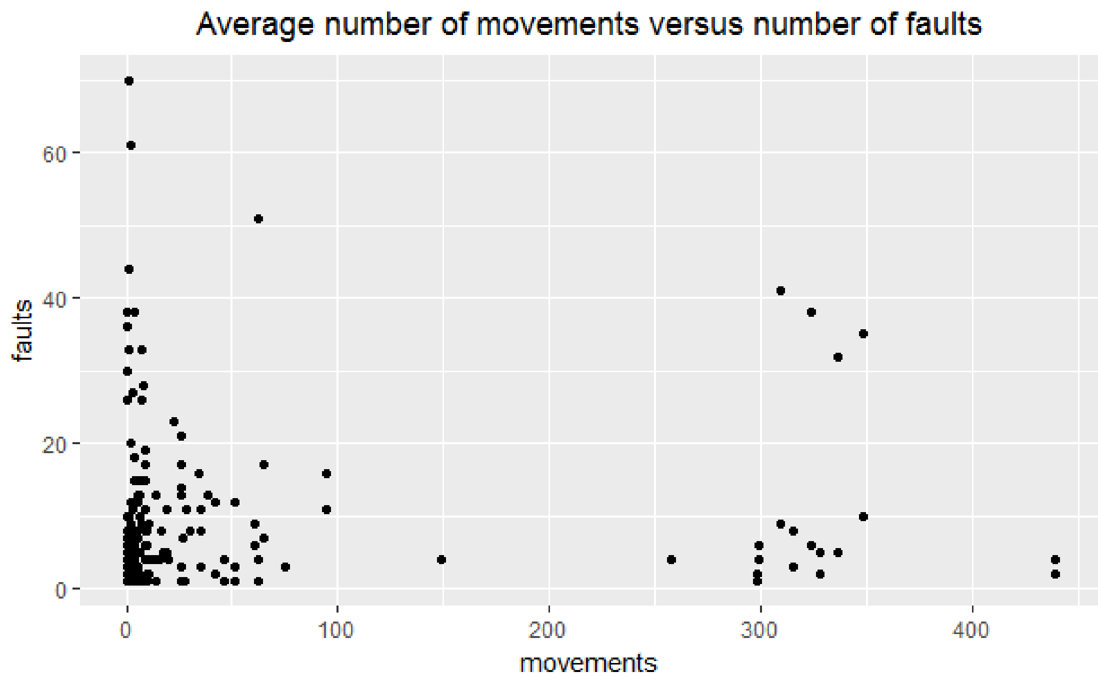

This plot reveals no particular relationship between number of movements and number of faults. The plot suggests there might be value in examing the data on a log-log scale. (A cluster of points near the origin with many points along both axes is characteristic of variables that may have a relationship on a log-log scale.) Here is the same data on a log transformed scale:

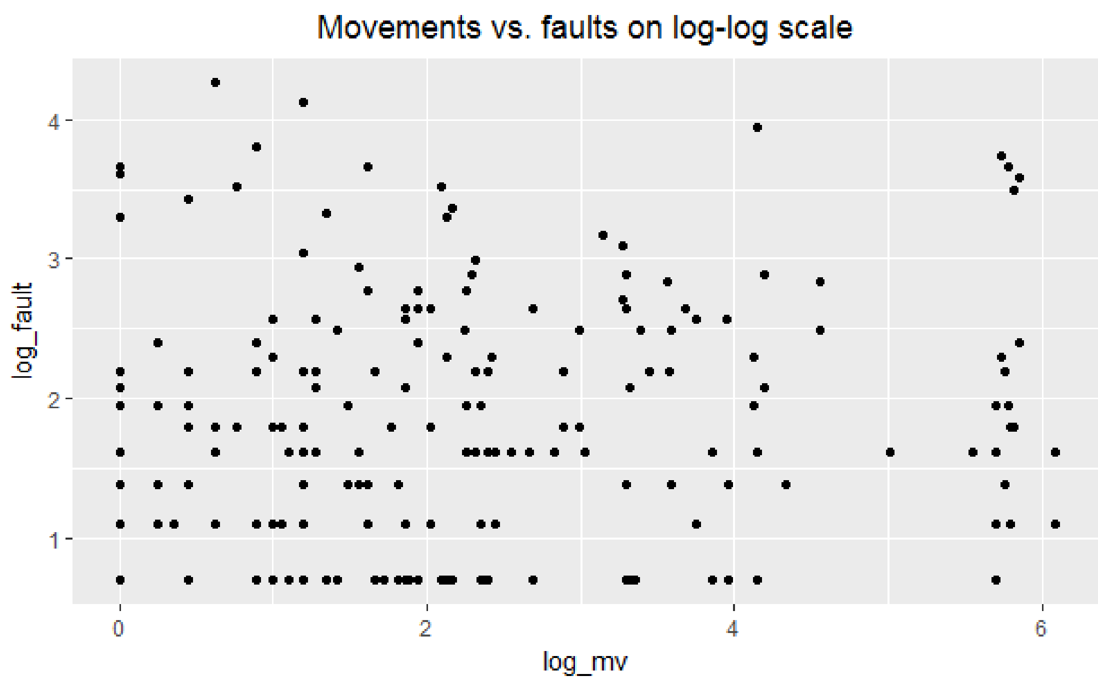

This plot confirms the lack of relationship between movements and faults.

It also may be deceiving to use the raw movement average, since many switches are seldom used and the movement data is only intended as a ranking mechanism more than a raw value. The following plot replaces the raw number of movements with their ranking within the data.

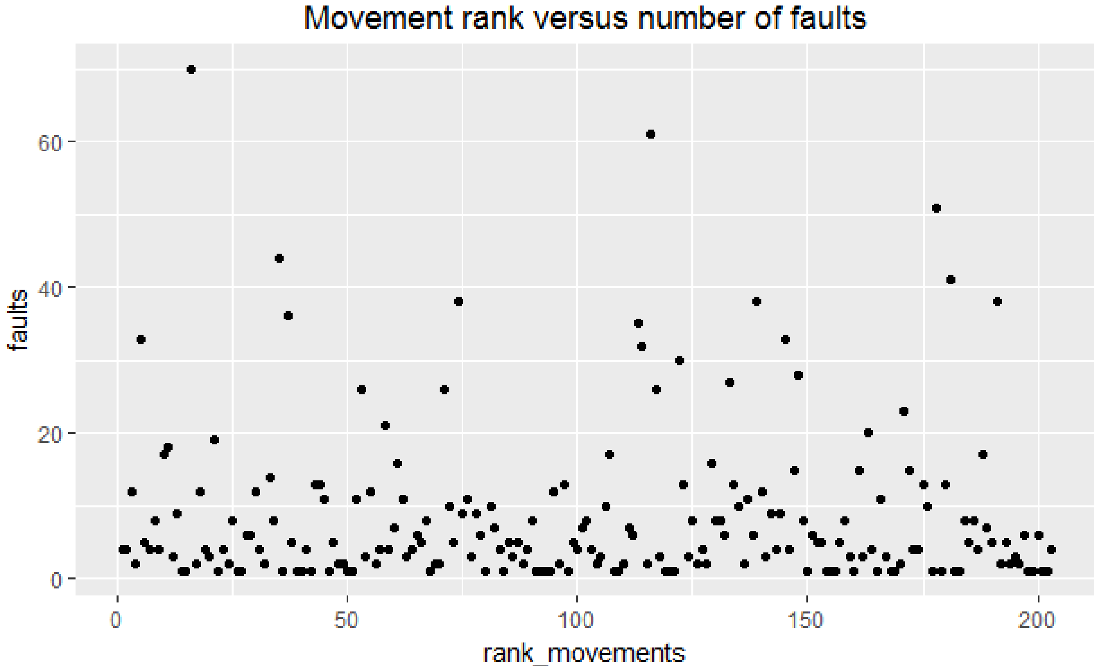

This plot provides more confirmation of the lack of relationship between movements and faults.

Controlling for the impact of external causes

It is possible that some switches are more prone than others to faults deemed "external" in nature. We performed a simple search on the terms: ice, debris, snow, and obstruction in the `HOW_REMEDIED` field of the faults data to identify those faults that were likely not internal to the switch itself.

We found it made no difference to the basic analysis after these probably external faults are excluded. The following plot shows excluding the external faults simply removes a small number of faults, roughly in proportion to the number of faults suffered.

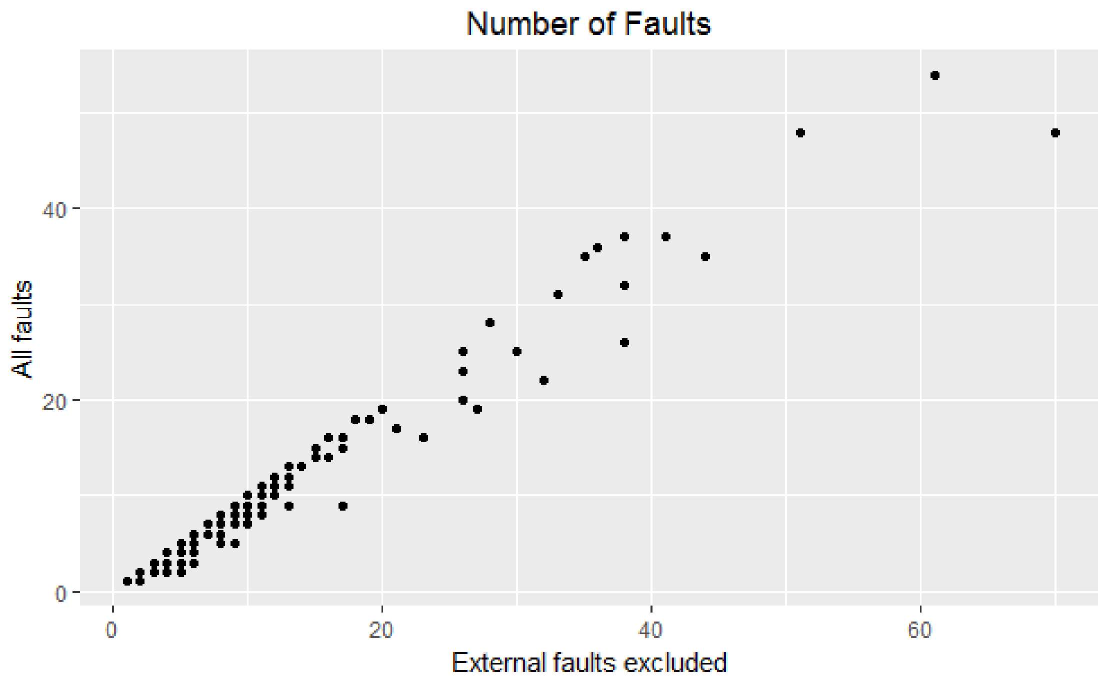

Impact of exposure to elements on faults

It is also possible that switches located outdoors could be more prone to faults. All yard switches were deemed to be outdoors. Otherwise the indoor/outdoor classification was provided by the TTC. It is not specifically recorded in the data.

There is a small impact of exposure on faults, which is mostly due to the faults classified as "external" due to ice, snow, debris, and obstructions.

The following plot shows a slight increase in number of faults for outdoor switches, all fault types considered:

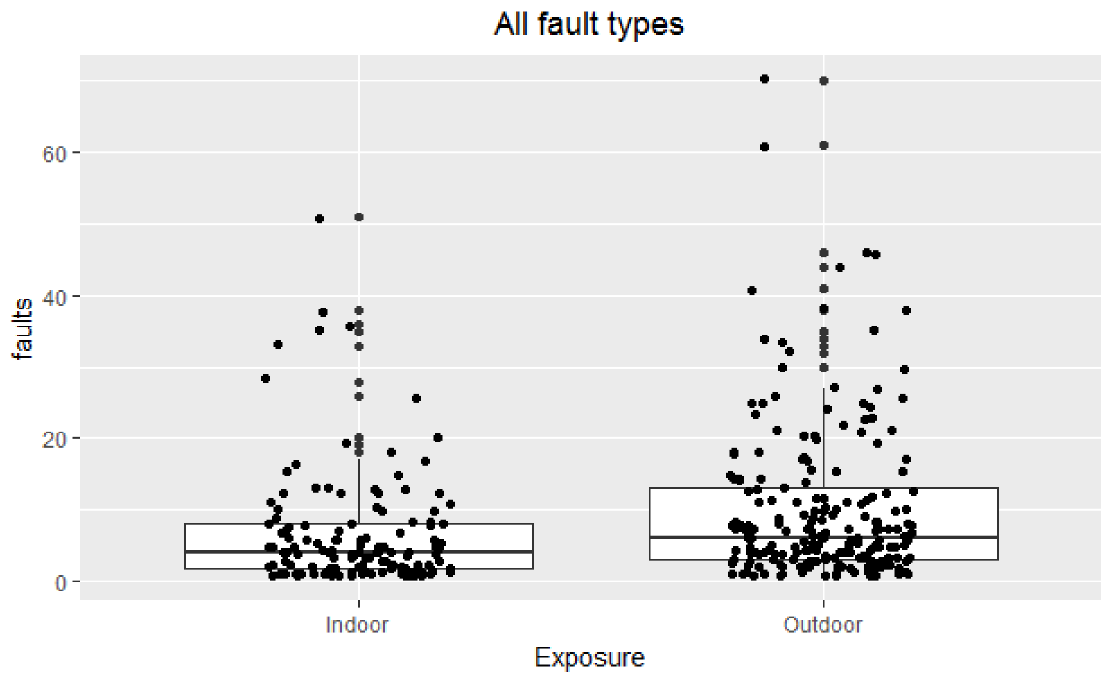

After removing the external faults there is a still a small increase for outdoor switches.

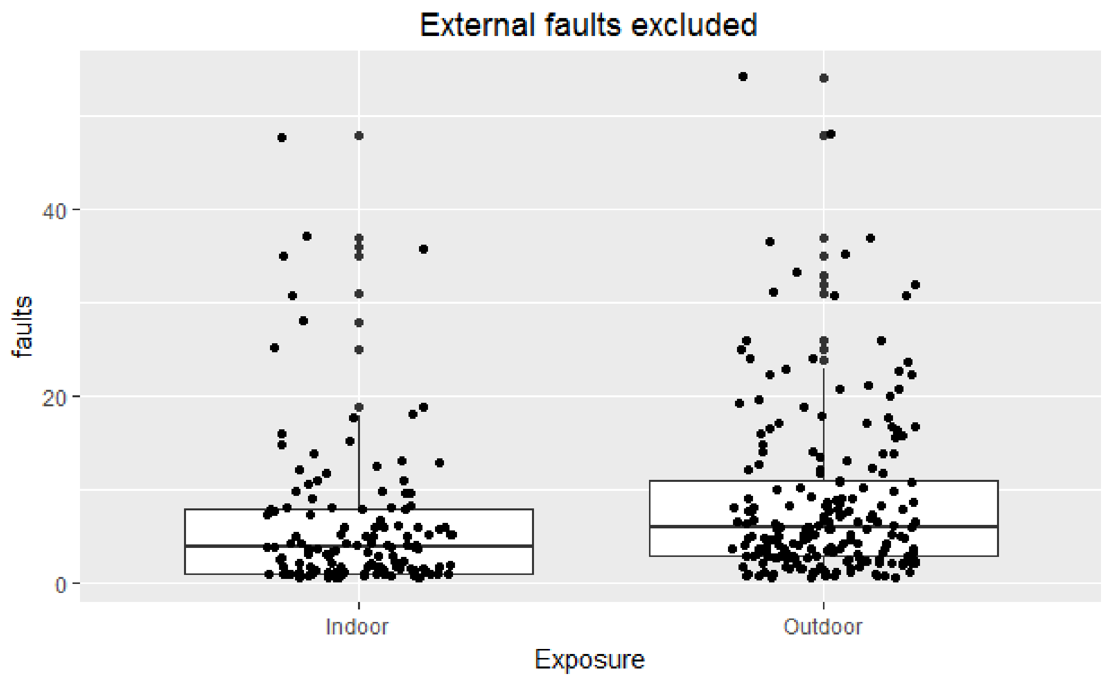

The same pattern holds when excluding yard switches. There is still a small increase in faults for outdoor switches. We could find no other explanation for this phenomenon.

Impact of age and renewal (very limited data)

There were 92 switch locations that have received one refurbishment, and 12 of these that have received two refurbishments. There are 24 known locations that have had a new switch installed. We will call these events "renewals".

There is no renewal information available for any of the other locations.

We examined the average number of faults per location per year, comparing before and after a renewal event took place. We found no evidence in the data of any impact of the renewal event, however the data are limited in several respects. We lack "after" data for most locations. Most switches have few faults, and after breaking down the data by switch and year of fault we are left with a small number of faults per year.

We do not feel this particular analysis can be used to inform any particular decision due to the nature of the data.

If per-switch age and renewal information were known, rather than only per-location partial information, it would be possible to more fully analyse the impact of switch age on number of faults.

Comparison of Terminal Stations

The terminal stations of each subway line (Line 1: Finch (FI) and Downsview (DN); Line 2: Kipling (KI) and Kennedy (KY); Line 4: Don Mills (DM) and Sheppard-Yonge (YN)) have switches that are used most often in the system, and the layout design for each pair is the same. We were therefore requested to analyze the data for these pairs of stations.

For Lines 1 and 4 there were only 36 and 37 total faults recorded for switch machines. For Line 2 there are 176 faults. For each line the fault numbers are roughly evenly distributed.

With so few faults, TTC staff could examinine them through an engineer's eye and learn more than a data analysis is likely to produce in this case.

Before/After January 2014

A basic method for assessing the overall, system-wide performance of a maintenance program is to examine calendar time versus cumulative numbers of faults, as in the following plot covering the entire time period of the data:

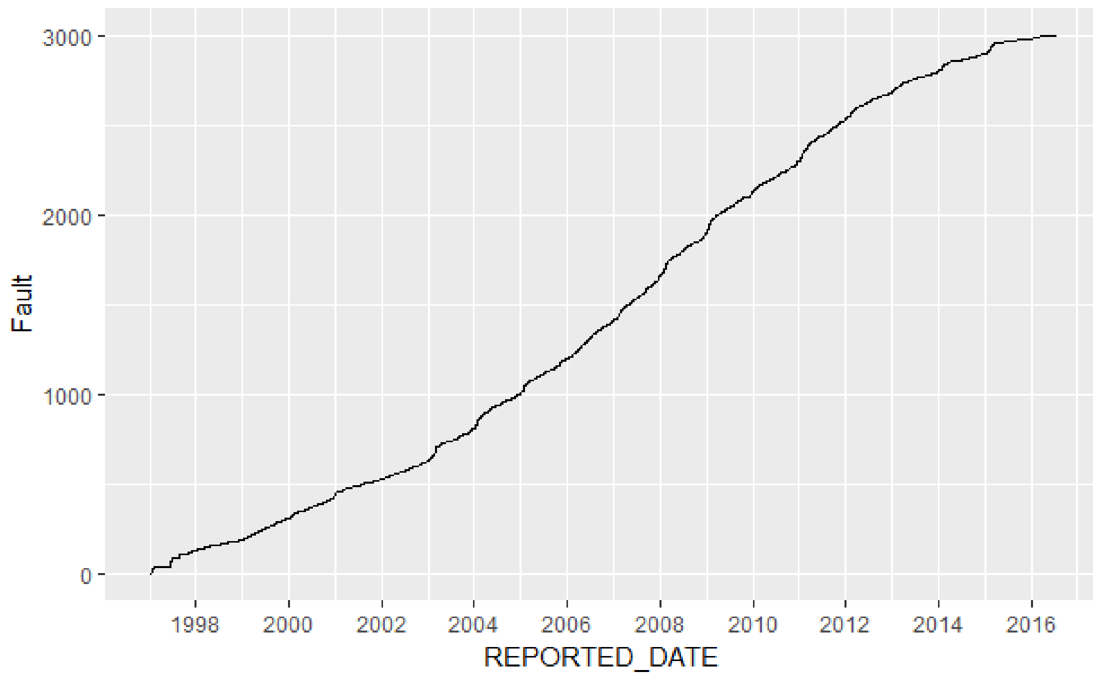

From this plot one can see a gradual change in overall reliability performance, which was deteriorating from 1998 to about 2007, and improving since that time up to the current date. There is no obvious performance change near the beginning of 2014.

By Manufacturer/Vendor

The following table shows substantial variation in the number of faults per switch machine when grouped by manufacturer. These differences may be explained by the unknown age factor, or some other systematic difference between the switch machines that is correlated with manufacturer.

MANUFACTURER	Faults	Count	Ratio
GEC	815	141	5.8
GRS	1859	161	11.5
US&S	78	19	4.1
WABCO	225	43	5.2
NA	21	1	21.0

Recommendations

Our main recommendations would be for switch machines to be tracked directly and individually in the asset management systems, rather than by reference to a location. This would also include determining installation dates. The lack of individual asset date precludes the use of many basic reliability and maintenance modeling techniques. We were restricted to overall, population-based analyses. This recommendation is not specific to switch machines. It would also apply to any asset class important to operations which is not already individually tracked.

Appendix - Analysis post-2007 excluding yard switch machines

I was asked to repeat the analyses using only non-yard switches and for faults reported after 2007-01-01.

Yard switches and mainline switches are distinguished by a column in the master list. The sample size is now 1108, down from 2998 in the originial analysis. We also excluded the decommissioned DN13CW, DN13DW, DN15CW, and DN15DW switches.

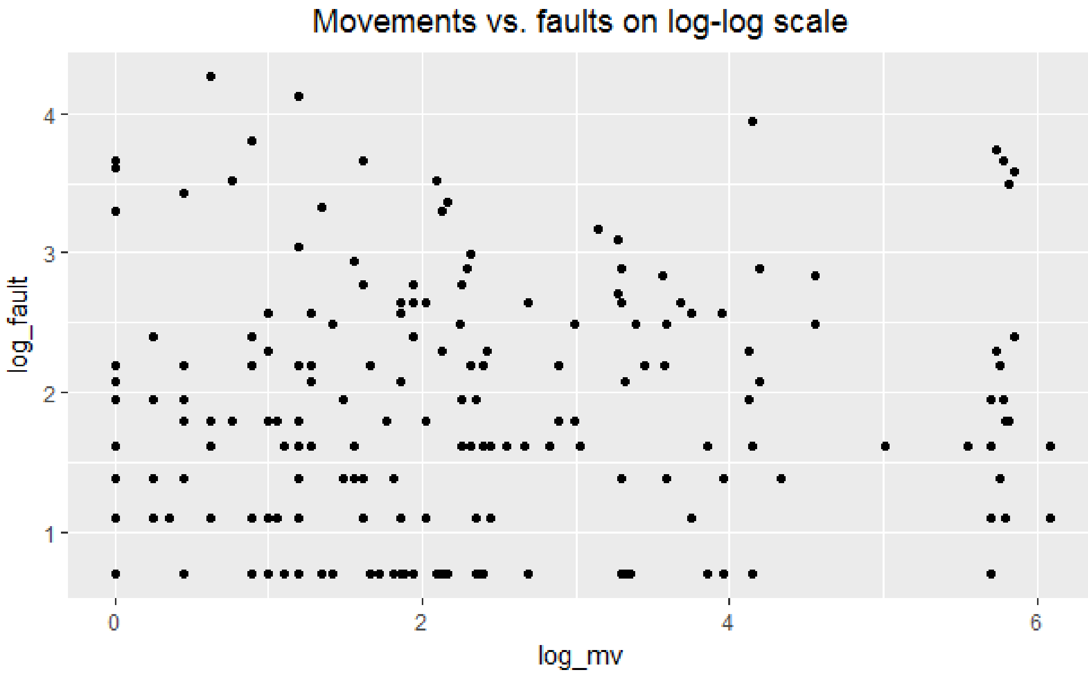

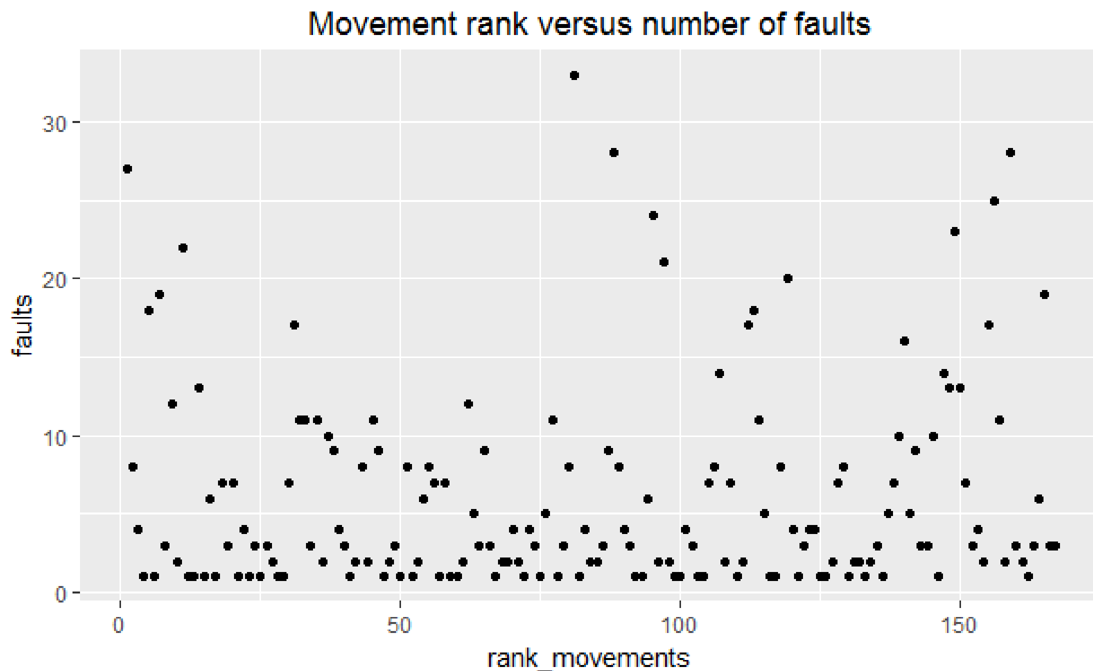
Movement rank versus number of faults
faults
rank_movements
0
10
20
30
0
50
100
150

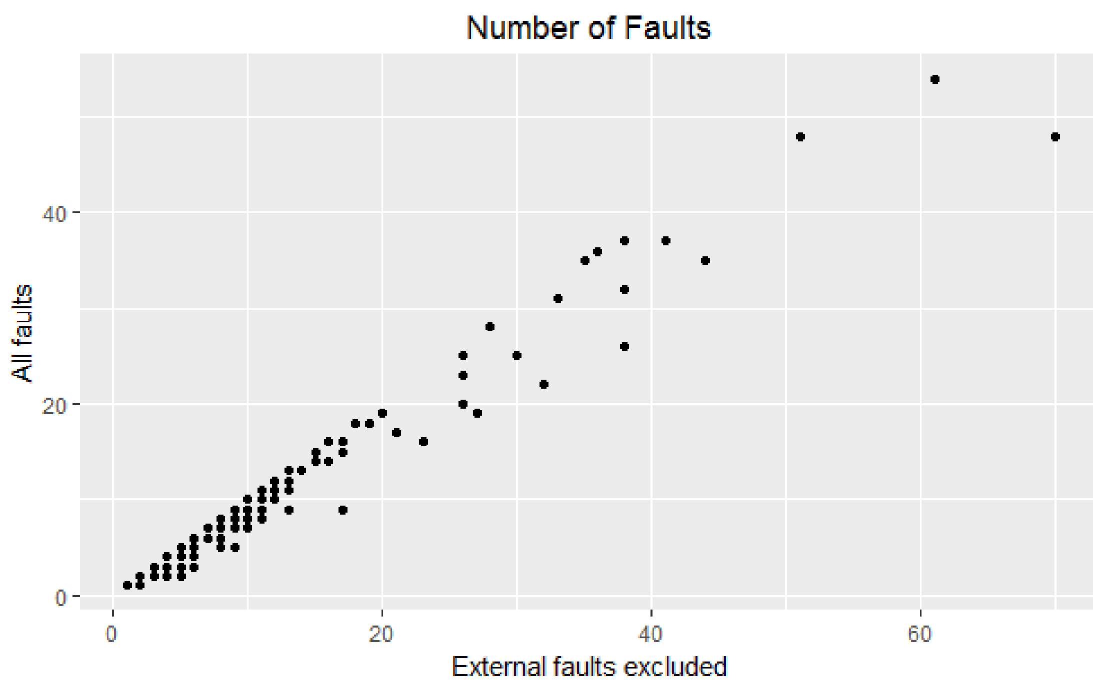
Number of Faults
All faults
External faults excluded
0
20
40
0
20
40
60

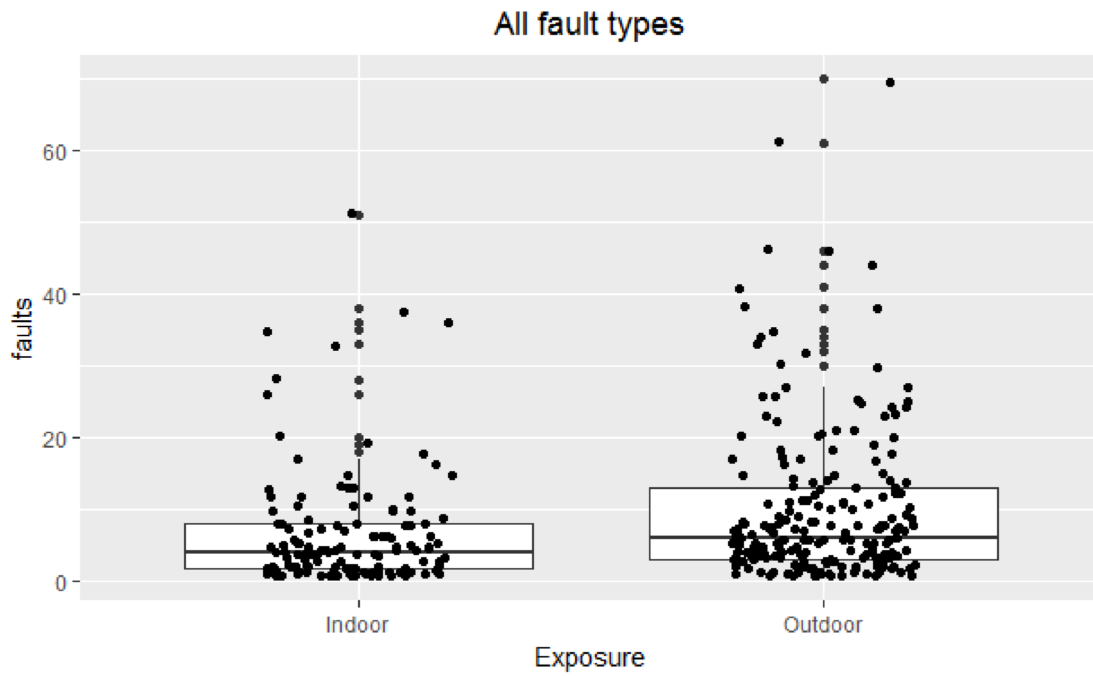

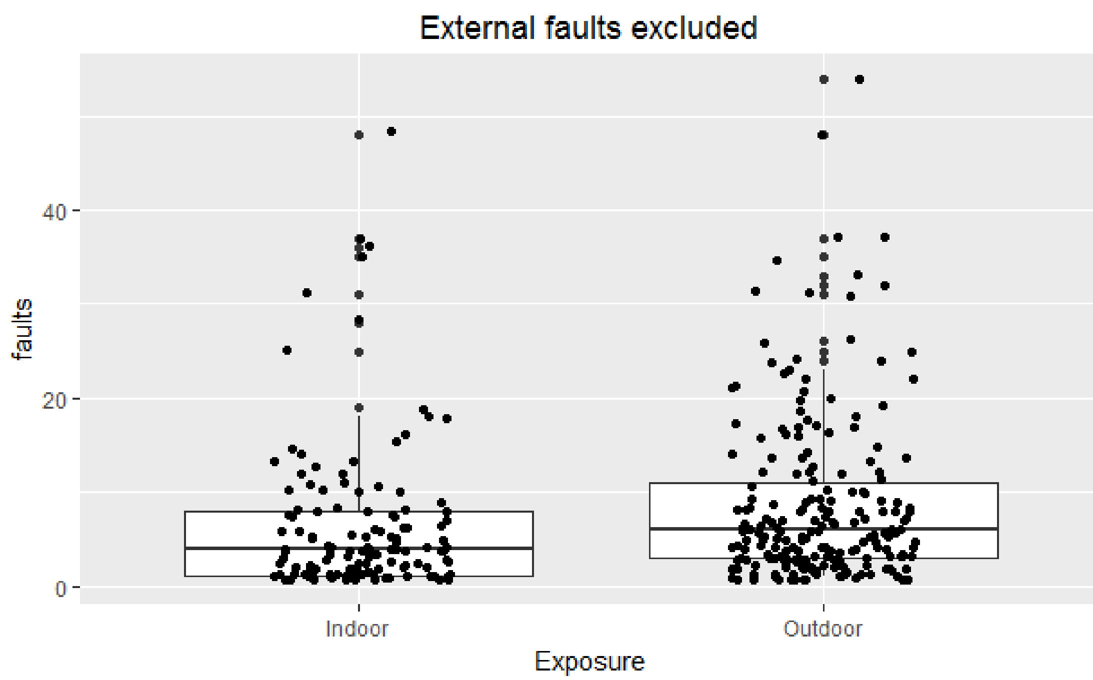

The following analysis also excludes all WABCO switch machines.

MANUFACTURER	Faults	Count	Ratio
GEC	490	126	3.9
GRS	571	84	6.8
US&S	40	8	5.0

TTC Trainstop Case Study

Neil Montgomery, C-MORE

Problem statement

Train stops are safety-critical automated devices at track level that engage train braking systems when trains violate the train blocking requirements (too close to the train ahead.) When faulty, go into a "fail safe" condition. Train stop faults are the leading cause of customer disruption in the subway system.

Data challenges

Most of the work on this project has been behind-the-scenes data preparation. Train stops are identified differently in the asset register than they are by front line staff interacting with them directly, who tend to use the location as assed ID. Eventually TTC had to manually generate a table with a correspondence between `ASSETNUM` and the location of the trainstop. Only just recently were some basic analyses possible.

Non-external faults by location and vibration status

The following failure modes (inferred by work done as in the `HOW_REMEDIED` variable) are deemed "external" to the trainstop and excluded in this basic analysis:

nff | debris | pop | snow | ice | gravel | ballast | balast | no fault found | all ok

These tables show counts and proportions of trainstops categorized by indoor/outdoor and high shock and vibration statuses.

For failure maintenance:

OUTDOOR / INDOOR	HIGH SHOCK & VIBRATION	Count	Prop	Fault Count	Fault Prop
Indoor	high s&v	47	0.04	300	0.07
Indoor	NA	843	0.79	2673	0.63
Outdoor	high s&v	17	0.02	156	0.04
Outdoor	NA	165	0.15	1094	0.26

For corrective maintenance:

OUTDOOR / INDOOR	HIGH SHOCK & VIBRATION	Count	Prop	CM Count	CM Prop
Indoor	high s&v	47	0.04	52	0.04
Indoor	NA	843	0.79	958	0.79
Outdoor	high s&v	17	0.02	21	0.02
Outdoor	NA	165	0.15	178	0.15

TTC Track Inspection Case Study

Neil Montgomery, C-MORE

Problem statement

The TTC performs visual inspections to monitor rail health of its subway system. The entire system is covered every 7 days. Additionally, non-destructive testing (NDT) is performed system-wise with a much smaller team. The entire system is covered every year. Incipient faults are re-inspected by the NDT team, and it is difficult for them to keep up with demand for inspections. Total rail failures (cracked) that actually occur in practice tend to arise between annual inspections, while the reinpected incipient faults tend not to progress.

We have been asked to determine if the inspection schedule(s) can be modified to prioritize areas of track by history, track type, or track geometry, while maintaining or improving reliability.

Track asset register files

We received a manually compiled asset register for all mainline tracks, with the following variables:

We added a new variable which we have tentatively called `Segment`, which concatenates one of `BD`, `YUS`, and `SHP` along with the row number, just to have an ID for each row in each spreadsheet of this workbook.

Each segment is physically identified by line direction and location (from a reference point in feet or meters).

Fault file

We also received an extract of fault records from the MOWIS system.

This file contains 211535 fault records. A fault is uniquely identified by the `DEFECT_NUMBER` variable. There are 166545 unique defects numbers. Defect numbers are repeated either when the defect is not immediately repaired and is re-inspected at some later time, or when a user of the database decides to make a new record to document the repair.

The location of the fault is given by the same marker and offset system as the asset register data.

BD line initial results

The YUS line has some challenges due to the way locations are physically identified, so we start with the BD line.

We can perform a merge of the fault data with the asset register so that each fault is associated to the track segment, section, geometry, and type.

High Priority Defects by Section and Geometry

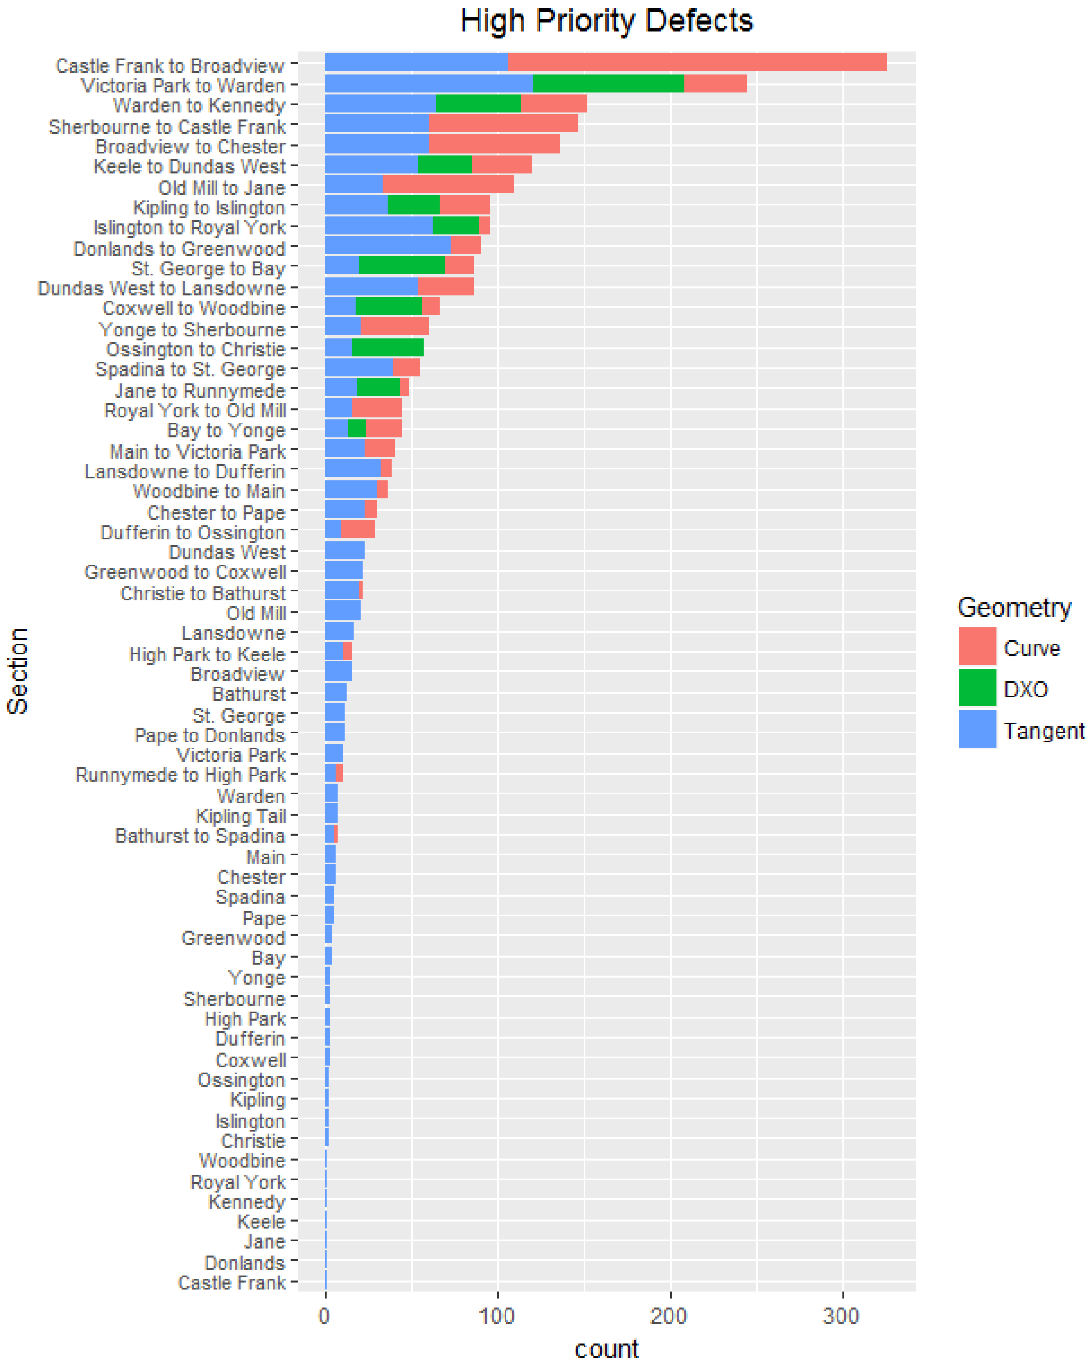

"Red" Defects by Section

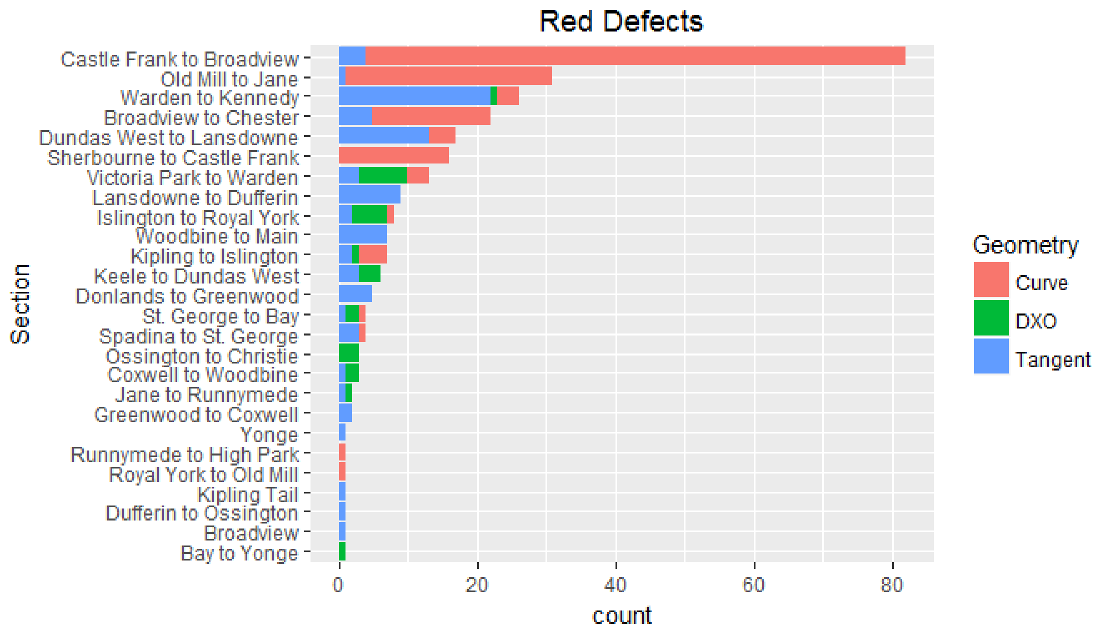

Table of HPD by geometry

Geometry	Count	Prop	Fault Count	Fault Prop
Curve	167	0.33	858	0.34
DXO	34	0.07	392	0.16
Tangent	300	0.60	1248	0.50

Faults by Geometry and Structure Type

Tracks can be constructed in a variety of ways. Abbrevations in the data are:

1. DCT – Double Concrete Ties
2. T&B – Wooden Ties and Ballast
3. DF – Direct Fixation of track to the concrete slab
4. STW – Special Track Works included the switches, diamond crossings
5. Guarded Curves – Curves with restraining rails or guard rails to prevent train derailment.
6. Unguarded Curves – without restraining rails or guard rails for high radius curves.

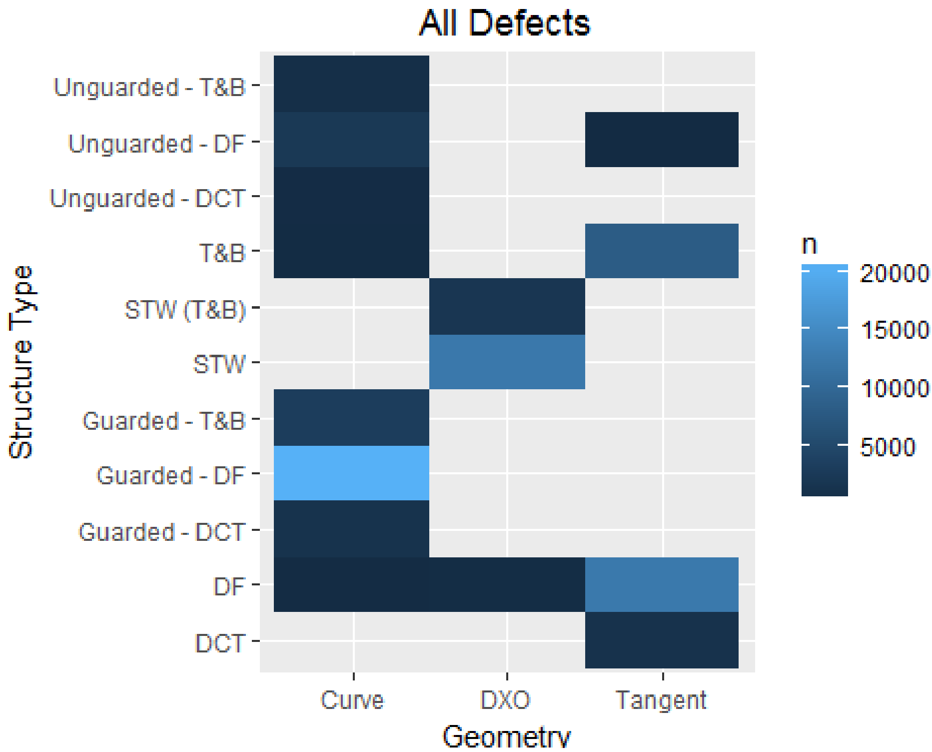
All Defects
Unguarded - T&B
Unguarded - DF
Unguarded - DCT
T&B
STW (T&B)
STW
Guarded - T&B
Guarded - DF
Guarded - DCT
DF
DCT
Structure Type
Curve
DXO
Tangent
Geometry
n
20000
15000
10000
5000

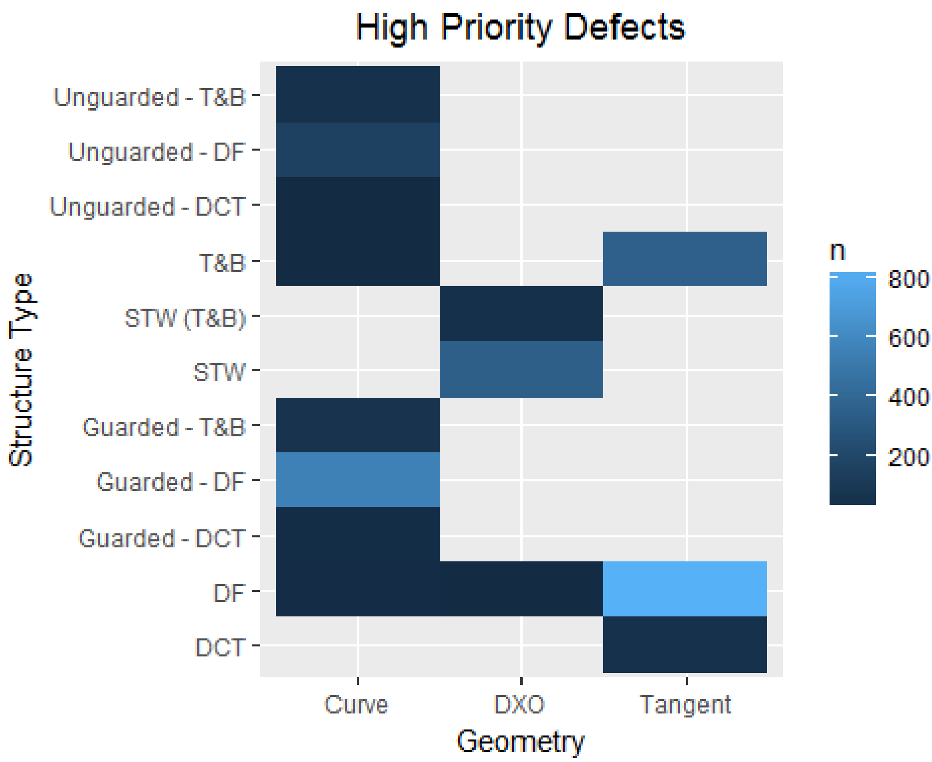
High Priority Defects
Unguarded - T&B
Unguarded - DF
Unguarded - DCT
T&B
STW (T&B)
STW
Guarded - T&B
Guarded - DF
Guarded - DCT
DF
DCT
Structure Type
Curve
DXO
Tangent
Geometry
n
800
600
400
200

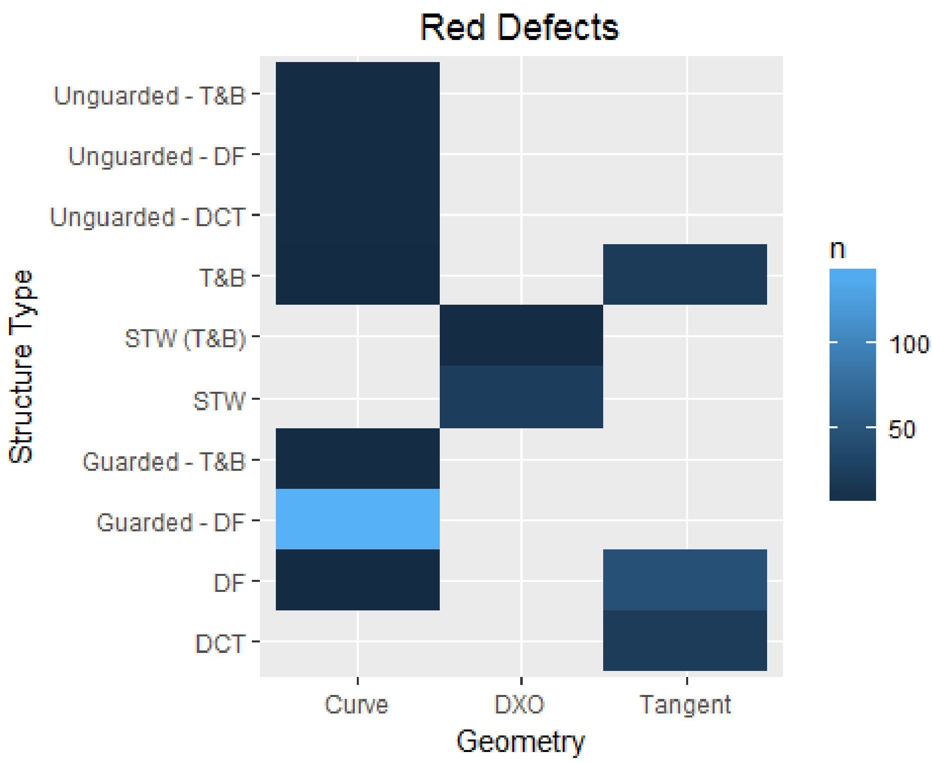
Red Defects
Structure Type
Unguarded - T&B
Unguarded - DF
Unguarded - DCT
T&B
STW (T&B)
STW
Guarded - T&B
Guarded - DF
DF
DCT
Curve
DXO
Tangent
Geometry
n
100
50

CB Criticality Tool

Neil Montgomery, Dragan Banjevic C-MORE

Arthur Sokol, CB

This document establishes methodologies for the main challenge of the project: extraction of Customer Basic Info. See *C-MORE / Canadian Bearings Criticality Tool Design and Scope V3* for context. We use the Process/Equipment/Component hierarchy in this document.

We will produce a separate document with a detailed *script* to use when meeting with customers.

The purpose of this extraction tool is to establish a list of "components" for which holding spares is a suitable strategy and that CB can supply, along with business impacts of components that can partially determine a criticality ranking.

For short, we'll call such a suitable component: "CB spareable".

Customer Basic Info could be extracted using a team from the customer consisting of mid- to high-level personnel along with front-line maintainers and operators. The extraction could take place during one or more sessions. Considerations:

- Will front line personnel be forthcoming in the presence of management?
- Would it me more effective to get a sense of overall business objectives first (from management), then focus on equipment and component details (from front lines)?

We call the latter strategy the "top-down approach", which we describe in some detail. A "bottom-up approach" could also be used, which we have also described. In brief the bottom-up approach uses the same overall structure as the top-down, but starting at a lower level and moving up and down as the company processes are explored.

Top-Down approach

Information can be extracted in the following structured manner. Wording, of course, to be discussed. At each step some assessment of important/impact is gathered that can help prioritize the next phase of questioning and can also produce an interim ranking for immediate self-assessment (to be combined with CB basic info to produce final criticality ranking.)

1. **Business focus:** In your words, what does your company do?
2. **Key impacts on business:** What are the keys to the business/production, and what are the keys to your safety and environmental concerns?
   2.1. Company may optionally be asked to rank the key areas.
3. **KPIs:** What are your KPIs in these areas, e.g. availability, "overall equipment effectiveness", production quota, no accidents/no incidents, etc.?
   3.1. **If company has KPIs – determine order of attack**

3.1.1. Rank the KPIs in order of importance or interest.
3.1.2. Which KPIs would you like to look at first?
3.1.3. How are you performing with respect to your KPIs?
3.1.3.1. Self-assess each KPI on a scale of 1 to 10.
3.1.3.2. If low, why? (e.g. perhaps completely unrelated to parts issues…)

Use company rank, or create rank if necessary.

3.2. **If company has no KPIs**
3.2.1. Either direct company towards a basic KPI such as availability.
3.2.2. Or just assume availability as a KPI and proceed to process information extraction.

4. **For each KPI in order of attack**

4.1. **Process list:** Which processes have the most impact on *this KPI?*
4.1.1. Are these all processes, or just the key processes?
4.1.2. Which process would you add, if you had to add just one more?

These two questions are "envelope-pushers" to ensure nothing is missed.

4.1.3. Can you rank the processes, or score each on 1-10 with respect to their impact on *this KPI*?

Use company rank, or create rank if necessary.

4.2. **For each process for *this KPI***
4.2.1. **Equipment list**: Which equipment has an impact on *this process?*
4.2.1.1. Are these all equipment, or just the key equipment?
4.2.1.2. Which equipment would you add, if you had to add just one more?

"Envelope-pushers"

4.2.1.3. Can you rank the equipment, or score each on 1-10 with respect to their impact on *this process*?

Use company rank, or create rank if necessary.

4.2.2. **For each equipment for *this process***
4.2.2.1. How is the equipment used: continuously, on demand, in emergency?
4.2.2.2. **Component list (if equipment is too complex to be CB spareable)**: Which components have an impact on *this equipment*?
4.2.2.2.1. Are these all components, or just the key components?
4.2.2.2.2. Which component would you add, if you had to add just one more?

"Envelope-pushers"

4.2.2.2.3. Can you rank for the components, or score each on 1-10 with respect to their impact on *this process*?

Use company rank, or create rank if necessary.

4.2.2.3. **For each component for *this equipment***
4.2.2.3.1. If not CB spareable, either skip or move to optional sub-component level if there are CB spareable sub-components
4.2.2.3.2. If CB spareable, proceed to some detailed reliability questions
4.2.2.3.2.1. What happens when this component fails? How much downtime? What does it cost?

4.2.2.3.2.2. Does the component failure cause the equipment fail, stop, or operate at a lower rating? What is the downtime? What does it cost?

4.2.2.3.2.3. What maintenance tactic is used currently (run to failure, preventative maintenance, condition-monitoring, etc.)

4.2.2.3.2.4. Do you keep spares for this component?

4.2.2.3.2.4.1. How did you determine your stocking policy for this component?

4.2.2.3.2.4.2. Where do you keep the spare components?

4.2.2.3.2.4.3. What do you do if there are no spares available?

4.2.2.3.2.5. Do you have trouble getting replacements?

4.2.2.3.2.6. How often do you have problems with this component, per day/month/year?

4.2.2.3.2.6.1. How many of these identical components are in use?

4.2.2.3.2.6.2. Is your assessment in this question fleet-wide or specific to this location where the component is being used.

4.2.2.3.2.7. How long do you expect it to last?

4.2.2.3.2.7.1. Can you provide a plausible upper and lower limit to this expectation?

4.2.2.3.2.8. What is its design life?

4.2.2.3.2.8.1. What do you do when it reaches design life? Use it anyway/life extension project/replace immediately/other? How long after design life might you wait?

4.2.2.3.2.9. Would you be OK with a non-OEM equivalent part?

5. **Wrap-up questions (in context of list that has been produced, possibly with a preliminary ranking)**

5.1. Do you keep any other components in stock? (If yes, determine where it fits into the hierarchy.)

5.2. How do you determine your overall stocking policy?

5.3. Do you have any dead stock?

5.3.1. Why?

5.3.2. Can any of it be scrapped or sold?

5.4. What are your holding costs? (to be expanded)

5.5. Have we missed anything really obvious?

5.5.1. Missing process, KPI, equipment, or component?

5.6. Is this preliminary ranking inconsistent with your experience? If so, what's clearly wrong about it?

Bottom-up approach

This approach could be suitable with companies either advanced enough to have a good parts list already, or have such a simple operation that going to the component level immediately is unlikely to result in important misses.

We need to discuss with CB how to decide when to use Top-Down and when to use Bottom-Up. Sales staff could have customer knowledge that would help inform this decision.

We have not fully determined how a bottom-up approach would (or could) work. Here are our ideas at the moment.

1. **Do you have a fairly complete list of components?**
   1.1. If the answer is no, and the operation is simple enough, proceed directly to Top-Down approach at the equipment level.
   1.2. If the answer is yes, we will classify components by equipment and process.
2. **What are the most important groups of components to your operation?**
   2.1. What equipments do they operate in?
      2.1.1. Can you rank the equipments or asses them on a 1-10 scale?
      2.1.2. Use this ranking to proceed to 4.2.2 above, working back down to the component level for each equipment.
   2.2. Then for each equipment,
      2.2.1. What processes do this equipments operate in? Proceed to 4.2 above, etc.

Customer Criticality / CB Information Combination

Neil Montgomery

Updated: May 12, 2017

Concept

We extract "operational criticality" from customers using a structured interview. We assess "unfilled demand criticality" using information combined from customers and CB data. We then combine these two to obtain an overall criticality, which can be used to prioritize parts strategies (which may include "number of spares" calculations.)

Operational Criticality

The customer's business will be assessed according to a Process/Equipment/Component/Sub-Component (P/E/C/S) hierarchy. Any additional level below Process is optional.

Through questioning, each P/E/C/S unit ends up with a score from 1 to 10, which is the importance of that P/E/C/S unit relative to that unit's parent.

The goal is to combine scores from various levels so that:

- we can fairly rank P/E/C/S units no matter which level of the hierarchy they are in.
- we can calculate absolute scores (with respect to the whole business) or scores going back only one or more levels in the P/E/C/S hierarchy.
- scores from higher levels will contribute more than scores from lower levels.

Notation

c_i : score for process i, relative to the whole business
c_{ij} : score for equipment j, relative to process i
c_{ijk} : score for component k, relative to equipment j
c_{ijkl} : score for component l, relative to equipment k
${}_1c_{ij}$: computed score for equipment j from process i, adjusted for the score of process i
${}_2c_{ijk}$: computed score for component k from equipment j from process i, adjusted for the scores of equipment j and process

Formulae

$$
\begin{aligned}
{}_1c_{ij} &= \sqrt{c_i\, c_{ij}^{c_i/10}} \\
{}_1c_{ijk} &= \sqrt{c_{ij}\, c_{ijk}^{c_{ij}/10}} \\
{}_2c_{ijk} &= \sqrt{c_i\, {}_1c_{ijk}^{c_i/10}} = \sqrt{c_i \left(\sqrt{c_{ij}\, c_{ijk}^{c_{ij}/10}}\right)^{c_i/10}}
\end{aligned}
$$

And so on. The final two scores are "absolute" for the whole business. A relative score going up only one level in the hierarchy would be, for example, ${}_1c_{ijk}$. This would rank all components within one process only, for example.

Call an absolute Operational Criticality C_O.

Unfilled Demand Criticality

From a combination of customer input and CB data, we determine, for each part in the Operational Criticality assessment, the following information (for example):

- Lead time to obtain part (in a time unit such as "days")
- Cost of unfilled demand, per day
- Demand rate (a fleet-wide measure) (1/MTBD)
- Holding cost per day (optional)

As a basic measure of unfilled demand, we could multiply: lead time, cost per day, and demand rate, to obtain an average cost per unfilled demand for that part.

These average costs would them be normalized (divided by the sum of all average costs) to obtain an Unfilled Demand Criticality. We would multiply by 10 to make it on the same order of magnitude as C_O. This value would be called C_D.

Overall Criticality

An overall measure of criticality would be a weighted combination of C_O and C_D, such as:

$$C_O^{W_O} C_D^{W_D}$$

where $W_O + W_D = 1$. This idea can be extended to more than two contributors to overall criticality.

The overall criticality ranking can then be used to prioritize parts to determine the appropriate strategy, which could be to calculate the appropriate number of spares to have in stock. But another maintenance strategy might also be more suitable.