Analysing an Industrial Safety Process through Process Mining: A Case Study

Anastasiia Pika[1], Arthur H.M. ter Hofstede[1], Robert K. Perrons[1], Georg Grossmann[2], Markus Stumptner[2], Jim Cooley[3]

Abstract. The application of data analytics has delivered significant value in a broad range of industries, but has frequently failed to bring about operational efficiencies or process safety improvements in asset-intensive sectors despite a profound increase in the volume of digital information that is stored by companies operating in this domain. Process safety management is a major concern of asset-intensive organisations, as the consequences of incidents can be catastrophic. In order to prevent incidents from occurring, organisations enforce prescriptive safety procedures frequently known as "Permit to Work" systems. These procedures can be thought of as a workflow consisting of a set of well-defined steps that have to be performed in a certain order to minimise safety risks. In this paper, we apply analytical tools from the area of process mining to shed light on how safety processes are actually executed. Process mining constitutes a relatively recent area of research at the intersection of data analytics and business process management (BPM) that is concerned with extracting insights from event logs that record process executions. We report the results of a process mining case study based on a Permit to Work system in an asset-intensive organisation in Australia.

1 Introduction

Organisations in asset-intensive domains invest a lot of resources in the management of safety risks (Sinelnikov et al., 2015; Podgórski, 2015), as the cost of safety incidents in these domains can be prohibitive. As a part of health and safety

[1] A. Pika (✉), A.H.M. ter Hofstede, R.K. Perrons
Queensland University of Technology, Brisbane, Australia
e-mail:{a.pika, a.terhofstede, robert.perrons}@qut.edu.au

[2] G. Grossmann, M. Stumptner
University of South Australia, Adelaide, Australia
e-mail: {georg.grossmann, markus.stumptner}@unisa.edu.au

[3] J. Cooley
Origin Energy, Brisbane, Australia
e-mail: jim.cooley@originenergy.com.au

management strategies, organisations design safety processes, e.g. Permit to Work systems (Flin et al., 2000). Such processes consist of steps which have to be followed by employees in a certain order to minimise safety risks. Safety processes are often complex and understanding their operational mechanics is crucial for risk identification and process improvement. Large volumes of data capturing the actual conduct of these processes are recorded; however, it is often a challenge to get actionable insights from this data (Podgórski, 2015; Sinelnikov et al., 2015).

In this paper we propose the application of process mining to derive insights from the execution of safety processes. Process mining is concerned with the analysis of process execution data (van der Aalst, 2016) and provides, among others, methods for discovering how processes are actually executed (process discovery), where deviations from expected process flows occur (process conformance) as well as techniques for process performance analysis and process comparison. Process mining algorithms take as an input an event log which contains information about process executions. Most process mining algorithms require that at least process instance identifiers, process activities and their timestamps are captured in the event log. Event logs may contain other information, for example whether an activity was started or completed, or information related to cost or location. Some information can be domain-specific, e.g. the requested loan amount for a mortgage application or the medicare number of a patient. Safety processes involve the conduct of activities in a certain order and when they are supported by information systems this history of work is typically recorded in an event log which can be analysed through the application of process mining techniques.

We present the results of a case study in which process mining was used to analyse the Permit to Work system in an Australian asset-intensive organisation. The Permit to Work system forms an integral part of the company's "Life Saving Rules", a set of clearly articulated requirements that all employees must follow to mitigate risks. Process compliance is considered to be critical in the company, and instances of non-compliance are treated as serious incidents in their own right. Our initial goal was general in nature, namely to investigate what insights about the process could be obtained through the application of process mining. We conducted data collection, interpretation, and analysis iteratively in collaboration with the process experts. During this engagement specific questions about the Permit to Work system emerged concerning process conformance, performance and evolution. These questions guided our subsequent analyses. The case study uncovered the process complexity, identified parts of the process where deviations from expected process flows occur as well as factors affecting process performance. Our findings highlighted the need for process simplification which could potentially make it easier to communicate the process, thus reducing mistakes in its conduct.

In the following sections, we first discuss related work, position our work and highlight our contributions (Section 2). We then describe the case study design, methods and findings (Section 3) as well as directions for future work (Section 4).

2 Related Work

Organisations have been using occupational health and safety management systems (OHS MSs) since the mid-eighties to minimise the number of health and safety (H&S) incidents (Podgórski, 2015). An OHS MS is "a set of institutionalized, interrelated, and interacting strategic H&S management practices designed to establish and achieve occupational safety and health goals and objectives" (Yorio et al., 2015). Permit to Work systems are often a part of OHS MSs (Flin et al., 2000). To measure and evaluate different aspects of OHS MSs companies use indicators: *lagging indicators* are used for performance measurement, *structural performance indicators* are used to evaluate "system compliance with a given specification form", and *operational performance indicators* are used to evaluate "effectiveness of internal system processes" (Podgórski, 2015). Structural and operational performance indicators are also referred to as *leading indicators* (Podgórski, 2015). Another purpose of leading indicators is "to flag potential problems" and "help uncover weaknesses in the organization's procedures or employee behaviour" (Sinelnikov et al., 2015). The numbers of leading indicators tracked by companies can be very high (up to a few hundred) resulting in large volumes of data being recorded (Podgórski, 2015). It is often a challenge to get insights from such data (Podgórski, 2015) and "turning data into action is a struggle for many organizations" (Sinelnikov et al., 2015).

Despite the challenges organisations face trying to make sense of safety-related data, Tan et al. (2016) argue that they "may be in a position to make considerable progress by applying 'Big Data' analytical tools" in the oil & gas industry; and Podgórski (2015) argues that companies should make better use of operational performance indicators to get "a picture of how processes operate at the shop-floor level". *However, it is not shown how these ideas can be operationalised*.

H&S incidents are often linked to a poor safety culture (Nordlöf et al., 2015) which is often evaluated using questionnaires (Flin et al., 2000; Podgórski, 2015). However, this approach has been criticised because employees' answers can be affected by "social expectations about their behaviour" and frequent involvement of all employees in such questionnaires is costly (Podgórski, 2015).

Several applications of Big Data for risk mitigation were reported by Tan et al. (2016): the use of hands-free checklists by workers assembling equipment in the oil & gas industry to reduce the chance of making mistakes; the use of built-in sensors on assets to monitor manufacturing assembly operations and reduce operational risks; and the use of Big Data to detect risks in supply chains.

We, on the other hand, propose the use of process mining to gain insights into the execution and performance of H&S processes. Process mining is a research area at the intersection of data analytics and BPM and is concerned with extracting insights from event logs that result from the conduct of business processes (van der Aalst, 2016). Risk-aware BPM provides methods for dealing with risks in processes (Suriadi et al., 2014) including methods for mining process risks from event logs, e.g., a recommendation system that supports risk-informed decisions (Conforti et al., 2015) and a method that helps to evaluate and predict overall process risk (Pika et al., 2016). *Process mining techniques helped to gain insights into process efficiency in organisations across different industries; however, to the best of our knowledge, they have not yet been applied to safety processes.*

3 Case Study

3.1 Analysis Approach

The case study was conducted in an Australian energy company. We analysed the company's Permit to Work system, a safety process consisting of activities which must be performed by employees to minimise work hazards. The process is an important part of the company's risk management strategy. Failure to follow the process may pose significant health and safety risks.

Data collection, interpretation, and analysis were performed iteratively in collaboration with some company personnel with knowledge of both the process and associated information. We refer to these personnel as process experts. At the beginning of the case study we did not focus on any specific questions; our goal was to learn what insights about the process we can uncover through the application of process mining. We refer to this part of the case study as *Stage I* and discuss it in Section 3.2. After we presented our initial findings and demonstrated different process mining capabilities to the process stakeholders, they were excited to learn more about specific aspects of the process. Our subsequent analyses were guided by their questions. We refer to this part of the case study as *Stage II* and discuss it in Section 3.3.

The process is supported by different information systems which record the process execution data. We worked with the process experts to collect this data from different sources, interpret it, and create an event log suitable for process mining algorithms (in XES[4] format). The resulting event log contained 43,061 cases (i.e. process instances) which were started and completed during the period of 28/03/2008-8/12/2016. The data was analysed using the open source process mining framework ProM[5] and the process analysis tool Disco[6].

[4] www.xes-standard.org

[5] www.promtools.org

[6] https://fluxicon.com/disco/

3.2 Stage I

3.2.1 Goals and Methods

At the beginning of the case study the company did not have any specific questions and was interested to see what they could learn about the Permit to Work process by analysing its data. During this stage we applied a range of process mining techniques to shed light on different aspects of the process. To learn how the process is executed in practice, we used the Disco miner, a process discovery tool which mines a fuzzy model from an event log (Günther & Rozinat, 2012). Process conformance analysis was used to identify where and how often actual process behaviour deviates from expected process flows. We applied an alignment-based process conformance technique (van der Aalst et al., 2012) which uses as input an event log and a normative process model (i.e. a model which describes expected process behaviour). The technique creates the so-called alignments by relating events in the log to elements of the model and highlights discrepancies between the prescribed and the observed process behaviour. We first created a normative process model based on the process documentation which was subsequently updated based on input from the process experts. An alignment-based performance analysis approach (van der Aalst et al., 2012) was used to identify parts of the process with the longest execution times. To check whether process behaviour and performance is changing, we used an extensible framework (Pika et al., 2017) which allows us to define measures of resource and process behaviour and to track their evolution over time.

3.2.2 Findings

Process discovery revealed that the Permit to Work system is a complex process with more than 2,300 process variants. Figure 1 demonstrates the discovered process model which represents all observed process behaviour and highlights the complexity of the process. A company representative was impressed by the insights revealed by the process discovery result. He noted that some process participants had often cited the process complexity as a leading cause of lost productivity and inadvertent process non-compliance, and that the discovered model provides evidence for the need to simplify the process.

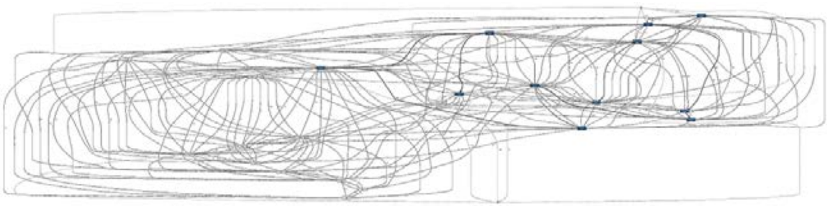

Figure 1: Discovered process model of the Permit to Work process.

Conformance checking identified those parts of the process where deviations from expected process flows occur. To further interpret the deviations, i.e. to learn whether they are related to skipped activities or activities executed in a wrong order, we checked those parts of the process where deviations occur using Disco. We discovered that many non-conformances occur at the beginning of the process and are related to risk assessment and approvals, for example, some approvals are skipped or in some cases authorisations are performed after a permit to work is withdrawn. We also learned that the frequency of some non-conformances is changing over time. Figure 2 shows an example of a process conformance issue (represented in terms of a fuzzy model) and its evolution over time. We can observe that in most cases (40,479) risk assessment was authorised (task *PTW_RA_Authorised*) after it was created (task *PTW_RA_Created*) which is expected process behaviour, while in 1,048 cases risk assessment was authorised before it was created which is unexpected. We can also observe that the frequency of such deviations is decreasing over time. The process experts could not explain this abnormal behaviour and further investigation is required to identify its causes.

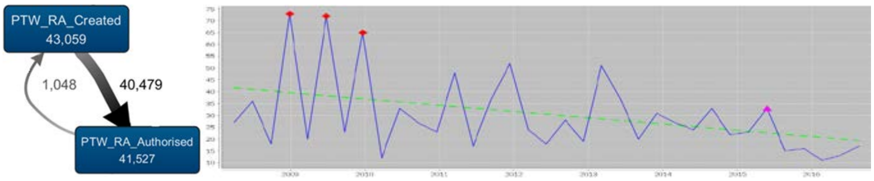

Figure 2: Example of a process conformance issue and its evolution over time: risk assessment was authorised before it was created.

Process performance analysis helped us to identify parts of the process and process variants with long execution times. For example, we learned that cases in which a permit was re-issued took on average 61.6 days while the mean duration of other cases was 33.8 days (mean durations of different variants of the process were checked using Disco). Similarly, cases in which testing was performed had longer execution times (mean duration 69.1 days) compared to other cases (mean duration 36 days). By checking process evolution, we learned that the process is getting more efficient over time. To get further insights into process performance, the process experts asked us to look at different types of work; we discuss these findings in Section 3.3.2.

3.3 Stage II

3.3.1 Goals and Methods

During the second stage of the case study we focused on specific areas of interest to the process stakeholders which emerged from Stage I. The stakeholders were interested to learn more about conformance issues of the process and their evolution over time. They asked us to analyse specific process behaviours related to Key Performance Indicators (KPIs) defined by the company. They were also interested to learn whether there are any differences in the execution and performance for different types of work. There is an attribute in the process execution data capturing type of work and the process experts grouped its values into four cohorts: *major outages*, *major forced outages*, *short-term maintenance* and *medium-term maintenance*. As per the company's request, during the second stage of the case study we analysed more recent data (cases started and completed during the period of 1/1/2011-8/12/2016).

In addition to process mining techniques described in Section 3.2.1, during the second stage of the case study we applied the ProM plug-in "Process Profiler 3D"[7]. This plug-in takes as input an event log and a process model and visualises differences in the execution and performance of different process cohorts. We used it to compare the aforementioned different types of work.

3.3.2 Findings

The process experts provided us with a list of process-related KPIs they were interested in. Here we show examples of findings for three KPIs: KPI 1 "Request for access must be approved seven days before the scheduled start date"; KPI 2 "Permit to work must be issued before the scheduled start date"; and KPI 3 "Permit to work must be withdrawn before the due date". Scheduled start date and due date are not process steps but they are recorded in the event log. The presence of this information allowed us to apply the same approach for checking process evolution (Pika et al., 2017) we used during the first stage of the case study. Figure 3 shows that the percentages of cases with KPI violations are different for the three KPIs: around 30% for KPI 1, 70% for KPI 2 and 50% for KPI 3 (on average across the whole period). Figure 3 also demonstrates that the level of violations does not change significantly for KPIs 1 and 2, but it is increasing for KPI 3. Overall, the number of non-conformances with respect to these KPIs is high; causes of this behaviour are yet to be investigated.

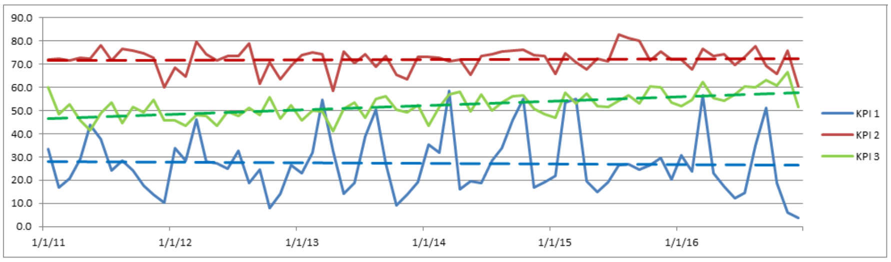

Figure 3: Percentage of cases started per month in which a KPI was violated

Process comparison identified differences in the execution and performance of the four types of work. For example, we learned that permit suspension is never repeated for *major forced outages*, but it can be repeated for other types of work.

---

[7] https://svn.win.tue.nl/repos/prom/Packages/ProcessProfiler3D

We discovered that there are significant differences in performance for the four cohorts, with mean case durations ranging from 7.5 days for *major forced outages* to 58.1 days for *major outages* (mean case durations were checked using Disco). Execution times also vary for the four cohorts in different parts of the process. Figure 4 shows mean waiting times for the four cohorts in a part of the process (arc colours represent different cohorts and arc heights reflect mean waiting times). While mean waiting time is the highest for *major outages* in most parts of the process (red arcs), this is not the case for mean waiting time between tasks *RFA_Created* and *RFA_Scheduled* which is the longest (7.7 days) for *medium-term maintenance* (yellow arc).

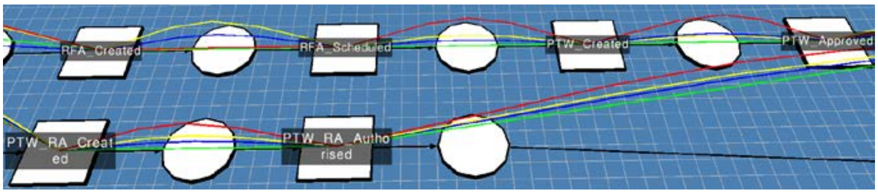

Figure 4: Mean waiting times for the four cohorts: major outages (red), major forced outages (green), short-term maintenance (blue) and medium-term maintenance (yellow).

4 Lessons Learned and Future Work

The application of process mining techniques allowed us to uncover the complexity of the permit to work process, identify those parts of the process where deviations from expected process behaviour occur as well as process variants with the longest execution times, discover differences in the execution and performance for different types of work, and learn that some process behaviour is changing over time. The company is considering an improvement initiative for the process aimed at simplifying it. This could potentially lead to a decrease in process conformance issues and to an improvement of process safety in the company.

The case study uncovered some data quality issues; for example, some process-related data is not recorded (e.g. activity start times) or is partially recorded (e.g. information about employees involved in the process). Accurate capture of all process execution data would enable a more elaborate analysis, e.g. we could consider activity durations or analyse behaviour of employees handling the process.

We also identified some shortcomings of alignment-based process conformance and performance analysis techniques. The conformance checking technique used can identify deviations from expected process behaviour, however further investigation is required to interpret these deviations. The performance analysis technique used can visualise execution times of different process parts; it could be enhanced by visualising performance of different process execution paths.

A step forward in the evolution of process safety is transforming the knowledge gained through the application of process mining into actions. Process mining helped us to identify some process conformance issues; however, we did not investigate the causes of these issues and ways to overcome them. As mentioned earlier, process complexity could be a factor affecting both process performance and conformance, and the company is considering a simplification initiative for the process. Lack of proper process communication or inconsistencies in process representations could be other possible sources of conformance issues. To improve process communication, it is necessary to capture the real world processes and depict how they relate to stakeholders and the affected artefacts, e.g. in the form of lifecycles (Schrefl and Stumptner, 2002). Inconsistencies in process representations can be detected (Grossmann et al., 2015) and corrected, e.g. with the help of process model repair (Polyvyanyy et al., 2016). Techniques for process configuration (van der Aalst et al., 2010; Mayer et al., 2011) can be used to support consistent representation of multiple process versions.

5 Conclusions

Safety processes aim to minimise safety hazards in asset-intensive organisations. A lot of data is recorded about the conduct of these processes, and whilst this data is extensively used to produce KPIs to measure business performance, it remains a challenge to obtain actionable insights about the processes themselves from this data. In this paper we reported the results of a case study in which process mining was applied to analyse process execution data of a Permit to Work system in an Australian energy company. The case study revealed different types of process insights, thus providing an initial indication of the potential of the application of process mining to safety processes in asset-intensive organisations. Specifically, the case study results highlight the complexity of the process and areas of poor process performance, providing evidence for the company for the need to simplify the process with a view to improving its reliability and efficiency.

6 References

Conforti, R., de Leoni, M., La Rosa, M., van der Aalst, W.M.P., & ter Hofstede, A.H.M. (2015). A recommendation system for predicting risks across multiple business process instances. *Decision Support Systems*, 69, 1-19.

Flin, R., Mearns, K., O'Connor, P., & Bryden, R. (2000). Measuring safety climate: identifying the common features. *Safety Science*, 34(1), 177-192.

Grossmann, G., Mafazi, S., Mayer, W., Schrefl, M., Stumptner, M. (2015). Change Propagation and Conflict Resolution for the Co-Evolution of Business Processes. *Int. J. Cooperative Inf. Syst.* 24(1).

Günther, C.W., & Rozinat, A. (2012). Disco: Discover Your Processes. In *Demonstration Track of the 10th International Conference on Business Process Management*, 40-44. http://ceur-ws.org/Vol-940/paper8.pdf

Mayer, W., Stumptner, M., Killisperger, P., Grossmann, G. (2011). A declarative framework for work process configuration. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 25(2), 143-162.

Nordlöf, H., Wiitavaara, B., Winblad, U., Wijk, K., & Westerling, R. (2015). Safety culture and reasons for risk-taking at a large steel-manufacturing company: Investigating the worker perspective. *Safety Science*, 73, 126-135.

Pika, A., Leyer, M., Wynn, M.T., Fidge, C.J., ter Hofstede, A.H.M. & van der Aalst, W.M.P. (2017). Mining resource profiles from event logs. *ACM Trans. Manage. Inf. Syst. 8, 1, Article 1* (in press).

Pika, A., van der Aalst, W.M.P., Wynn, M.T., Fidge, C.J., & ter Hofstede, A.H.M. (2016). Evaluating and predicting overall process risk using event logs. *Information Sciences*, 352, 98-120.

Podgórski, D. (2015). Measuring operational performance of OSH management system–A demonstration of AHP-based selection of leading key performance indicators. *Safety Science*, 73, 146-166.

Polyvyanyy, A., van der Aalst, W.M.P., ter Hofstede, A.H.M., & Wynn, M.T. (2016). Impact-driven process model repair. *ACM Trans. Softw. Eng. Methodol.*, 25(4), 28.

Schrefl, M., Stumptner, M. (2002). Behavior-consistent specialization of object life cycles. *ACM Trans. Softw. Eng. Methodol.* 11(1), 92-148.

Sinelnikov, S., Inouye, J., & Kerper, S. (2015). Using leading indicators to measure occupational health and safety performance. *Safety Science*, 72, 240-248.

Suriadi, S., Weiss, B., Winkelmann, A., ter Hofstede, A.H.M., Adams, M., Conforti, R., Fidge, C., La Rosa, M., Ouyang, C., Pika, A., Rosemann, M., and Wynn, M. (2014). Current research in risk-aware business process management – overview, comparison, and gap analysis. In *Communications of the Association for Information Systems*, volume 34, article 52, 2014.

Tan, K.H., Ortiz-Gallardo, V.G., & Perrons, R.K. (2016). Using Big Data to manage safety-related risk in the upstream oil & gas industry: A research agenda. *Energy Exploration & Exploitation*, 34(2), 282-289.

van der Aalst, W.M.P. (2016). *Process Mining: Data Science in Action:* Springer-Verlag, Berlin. http://www.springer.com/978-3-662-49850-7

van der Aalst, W.M.P., Adriansyah, A., & van Dongen, B. (2012). Replaying history on process models for conformance checking and performance analysis. *Wiley Interdisc. Reviews: Data Mining and Knowledge Discovery*, 2(2), 182-192.

van der Aalst, W.M.P., Dumas, M., Gottschalk, F., ter Hofstede, A.H.M., La Rosa, M., Mendling, J. (2010). Preserving correctness during business process model configuration. *Formal Asp. Comput.* 22(3-4): 459-482.

Yorio, P.L., Willmer, D.R., & Moore, S.M. (2015). Health and safety management systems through a multilevel and strategic management perspective: Theoretical and empirical considerations. *Safety Science*, 72, 221-228.