Precursors and startling lessons: Statistical analysis of 1250 events with safety significance from the civil nuclear sector

Journal Article**Author(s):**

Ayoub, Ali; Stankovski, Andrej; Kröger, Wolfgang; Sornette, Didier

Publication date:

2021-10

Permanent link:

<https://doi.org/10.3929/ethz-b-000490250>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

Reliability Engineering & System Safety 214, <https://doi.org/10.1016/j.ress.2021.107820>

Precursors and startling lessons: Statistical analysis of 1250 events with safety significance from the civil nuclear sector

Ali Ayoub^{a,b,1,*}, Andrej Stankovski^{a,1}, Wolfgang Kröger^{a,b}, Didier Sornette^{a,b,c}

^a Chair of Entrepreneurial Risks, D-MTEC, ETH Zürich, Switzerland

^b ETH Risk Center, ETH Zürich, Switzerland

^c Institute of Risk Analysis, Prediction and Management (Risks-X), Southern University of Science and Technology (SUSTech), Shenzhen, China

ARTICLE INFO

Keywords:

Nuclear accidents
Nuclear safety
Precursors
Operational experience
PSA
Risk
Statistical analysis

ABSTRACT

We analyze the ETH Zurich open curated database of 1250 worldwide nuclear events focused on safety significance with potentials for precursors, presented in the companion paper. We find that major accidents always trigger a wave of “reactive” reporting as well as changes in regulatory or corporate management that last 5 to 6 years, mostly due to increased alertness, improved transparency, uncovering latent design errors, and heightened public pressure. The leading causes for multi-unit events are found to be external triggers and design issues, confirming the need to adapt PSAs to cover multi-unit events accordingly. Common-cause failures (CCF) are found to occur fairly frequently, at different levels, and can significantly erode the safety of the plant. From the lessons learned from this analysis, we suggest that frequent review of components design and operating procedures, employing different teams for testing and maintenance activities on redundant trains, and sharing operational experience between plants of similar designs, are some of the steps that should be taken in order to limit future occurrences of CCFs and beyond that further improve plant safety. We identify some quantitative signs of aging for plants after the age of 25. Our findings stress the need for larger recording, reliance, and sharing of operational data to support learning from experience and avoid reoccurrence of accidents and events.

1. Introduction

Evolving large-scale engineered systems that our societies depend on, and which are inherently hazardous, call for predictive analysis of potential risks. As statistical information of direct use is limited by nature, analytical methods incorporating data and “lessons learned from the past” are needed to model the behavior of systems. The latter is based on compiling and evaluating data of experienced failures and events at different levels of abstraction. When dealing with interdependent critical infrastructure systems and their involved complexity, traditional risk modeling usually turns out to be insufficient to cover all their aspects and mechanisms [1]. A vivid example of less complex, but complicated, systems are nuclear power plants, made of thousands of components, systems, and structures, connected through diligent mechanisms. The corresponding high levels of redundancy and diversification, as well as separation, ensure utmost safety for their operation. Despite their high level of sophistication, nuclear power plants are not autonomous, as they require uninterrupted human surveillance,

operation, and maintenance. Additionally, a nuclear power plant is not an independent entity that can be modeled in isolation from its site and operational environment. On the one hand, a nuclear plant heavily relies on its external environment during its normal operation (offsite power grid, water intake). On the other hand, the external environment can be an extremely perturbing factor for the safe operation of a plant (earthquakes, floods, grid disturbances, power losses, water conditions). All these inherent complications, dependencies, man-machine interactions, external factors, and above all, the cultural and organizational aspects that shape many actions and consequences, make it extremely hard for mathematical models – especially ones aiming at assessing risks – to capture adequately the behavior of the plant. Nevertheless, traditional methods based on decomposition and causality, as well as human reliability analysis, appear to remain very useful and well suited.

As an illustrative example, probabilistic safety assessment (PSA) is a prominent risk framework that is extensively used in the nuclear power industry to identify design weaknesses and check whether the safety concept is balanced [2]. In particular, PSA models aim at quantifying

* Corresponding author.

E-mail address: aayoub@ethz.ch (A. Ayoub).

¹ These authors had equal contribution to this work.

nuclear power plants risks by:

- 1- postulating and analyzing what could go wrong from within the plant and from the outside;
- 2- how likely is it to happen, and if it happens;
- 3- what are the consequences.

With more than 40 years of utilization, PSAs have reached high levels of advancement that make them very robust and reliable in assessing risk. PSAs have even become part of the licensing process. However, due to the aforementioned complexities in nuclear power plants, PSAs still suffer major limitations, with completeness being a very prominent one [3,4], i.e. it is mathematically impossible to prove that PSA models can cover all conceivable accident scenarios and scopes.

Operational experience appears to be very useful in these situations, as it can provide a rich information source supporting risk models [5]. In addition to helping in the initial development of the models, lessons learned from operational experience often trigger amendments and improvements to the models. Besides, operational experience provides an empirical basis for the estimation of components reliability parameters and initiating events frequencies [6]. Furthermore, as many intrinsic features and system characteristics manifest themselves in empirical events, operational experience can be very useful to assess and validate mathematical models. Apart from modeling support, operational experience has been used to aid improvements and post-design modifications. For example, major accidents in the nuclear power industry have triggered industry-wide retrofits and stress tests [7], with the most prominent being the post TMI control room changes, post Chernobyl RBMK design changes, and the post Fukushima “Diverse and Flexible Coping Strategies FLEX” backfits [8].

Realizing the potential stemming from operational data, industries and researchers started paying increased attention to observe, record, and utilize real incidents and inadequacies to learn from mistakes and prevent their reoccurrence [9–12]. In this regards, Moura et al. [13] have compiled a dataset of 238 major accidents from different heavy industries, and developed a new approach to test and validate risk models, study causal interactions, and improve systems designs and robustness. Zhang and Mahadevan [5] utilized a database of historical aviation accidents and developed a Bayesian network to model causal relationships and facilitate probabilistic inference to support accidents investigation and risk analysis. Wheatley et al. [14] utilized a dataset of about 200 nuclear accidents and incidents, with several severity annotations, to empirically study the safety performance of nuclear power plants, and estimate the financial risks and likelihoods of major accidents.

Worldwide civil nuclear operations have accumulated more than 18'000 reactor-years [15] of experience, providing substantial amounts of data, a subset of which we try to utilize in this paper. Specifically, we are using the ETHZ Curated Nuclear Events Database [16,17], to extract important lessons and identify potential precursory signals. The database currently contains more than 1250 safety important events from around the world, making it by far the largest open-source, comprehensive, and academically driven initiative in the world. It serves all communities, from nuclear operators, regulatory agencies, scientists, and the public. It is used to provide various input for precursor analysis and other statistical analyses [18,19], with the aim to properly assess the safety of modern nuclear power plants, and identify trends and patterns of associated risks. With this amount of such high quality data, supported by its detailed and rich analyses and features, we aim at extracting reliable statistics leading to insights that can help averting future disasters or costly precursors.

2. Data structure and description

Operational data in the civil nuclear industry is well recorded at national levels, mandated by regulatory requirements (a vivid example

are the USNRC licensee event reports [20] and Accident Sequence Precursor program (ASP) [21,22]. At the international level, different efforts exist – with different scopes and aims – to collect operational experience; major examples include: World Association of Nuclear Operators [23], International Atomic Energy Agency Reporting System [24], European Clearinghouse on Operating Experience Feedback for Nuclear Power Plants [25]. However, these efforts in general are either not publicly available, actor-oriented, lack homogeneity (different reporting styles), and lack practical annotations and searchability (usually in the form of static reports).

Realizing these limitations, we have compiled a unique database, namely, the ETHZ Curated Nuclear Events Database [16], with more than 1250 safety significant events around the world, covering the whole time-period of civil nuclear power operation. The events in the database have been systematically analyzed by multiple nuclear-safety experts, providing a curated intermediate-level description, and coherent breakdown of features such as origin of the event, causes and contributors, type of failure, operating mode, failure sequence, significance, and many others. A detailed explanation of the database features, classification criteria, and definitions, can be found in the companion paper [26]. The database along with its custom-made graphical user-interface, provide an open access, user-friendly, and information-rich asset for the scientific community, industrial analysts, as well as the public.

3. Statistical analyses

The analyses, statistics, and interpretations presented in this paper are based on the data contained in our database. They are related to several different topics that we believe might give significant and novel insights for researchers and the industry, including the occurrence rate of events, the analysis of single and multi-unit sites, common cause failures (CCFs), quantitative aging signs, and the reliability of safety functions and systems.

3.1. Occurrence rate of events over time

In order to study the likelihood of occurrence of accidents and precursors, we examine the evolution of their rate over time. We define λ_t , the rate of occurrence of events per reactor in year t . Given that events are rare and are sparsely spread in time, we take the simplest assumption that their temporal distribution is Poissonian, in other words there is no correlation between successive events, and the distribution of waiting times between one event and the next one is exponential with rate λ_t . Then, the statistical estimate of λ_t is simply $\hat{\lambda}_t = \frac{N_t}{r_t}$, where N_t is the number of observed events in year t , and r_t is the number of operating nuclear reactors in year t . Under this assumption, the variance of λ_t is $Var(\lambda_t) = \frac{\lambda_t}{r_t}$, which can be calculated using the estimate $\hat{\lambda}_t$ to evaluate confidence intervals in a standard way. Fig. 1 shows the calculated rates, $\hat{\lambda}_t$ (solid dots), along with their Poisson standard errors (represented by the error bars) defined as the square-root of $Var(\lambda_t)$. Furthermore, the number of operating reactors over time, r_t , is also shown in Fig. 1 (continuous solid line).

It should be noted that only reactors which have event entries in the database are taken into account in this analysis in order to maintain consistency.

Four distinctive periods can be observed based on the occurrence rate of events:

- Early period (1965–1985) – early days of nuclear power defined by limited knowledge, insignificant operating experience, and lack of transparency. This period has a pronounced volatile nature with large variances in the rates of events, mainly due to the low reporting and low number of operating reactors. The average rate of events per reactor year, λ , of that period is quite high with $\lambda = 0.16$. During this

Fig. 1. Occurrence rate of events per reactor-year (solid dots with error bars, left axis) and the number of operating reactors appearing in our database (continuous solid line, right axis).

period, the Three Mile Island accident occurred (March 1979), which triggered a swarm of events afterwards, most probably due to increased safety awareness and design changes.

- Chernobyl cluster (1986–1992) – sudden increase in reported events triggered by the Chernobyl accident, with $\lambda = 0.13$ events per reactor year. This period is characterized by reactive reporting, increase of transparency, safety reviews and major changes, general negative public perception of nuclear power, as well as related public pressure on policy makers.
- Calm periods (1993 – 2010; 2017 – 2020) – largely “uneventful” periods where the implementation of lessons learned from the previous accidents, frequent inspections and the slightly diminishing public pressure might have led to a decrease in witnessed and reported events. A decline of attentiveness by plant operators and regulatory bodies might also contribute to this trend. The average rate of events per reactor year for these two periods is $\lambda = 0.07$ and $\lambda = 0.06$, respectively.
- Fukushima cluster (2011–2016) – in the wake of the Fukushima Daiichi nuclear accident, new waves of reactive reporting, uncovering of latent design deficits and vulnerabilities (e.g. by stress-tests), and regulatory changes were triggered, leading to a new spike in events with a $\lambda = 0.14$ events per reactor-year.

The effects of the Chernobyl and Fukushima accidents are evident, with pronounced periods of heightened awareness (characterized by an increase of events rates) lasting for 5 to 6 years. As discussed, a possible explanation for this trend appears to be the increase in transparency, public pressure and sudden uncovering of latent errors. A contributing factor could also be the typical duration of executives’ terms and mandates in office of 5–6 years, who tend to enforce a strict safety culture following major accidents, which in turn trickles down to employees and subordinates and allow for increased – intra and inter – reporting habits and transparency [27,28]. Following this 5–6-year period of alertness, we observe a calm period of low event rates, likely due to the effects of implemented safety and regulatory improvements. A concerning factor possibly contributing to this trend can be “lowering of the guard” of operators and regulators, on account of the diminishing public pressure over time. Therefore, in order to avoid another major accident after a calm period, licensees must retain good communication and transparency habits both internally and with the regulators. Similar to war drills performed by militaries to maintain combat readiness in times of peace, top management and regulators must demand a high level of attentiveness during “calm periods”. Furthermore, the plant and

operators’ response should constantly be challenged with postulated accident scenarios and simulations, in analogy to war drills. The regulators must also have an overview of whether all changes stemming from lessons learned are properly and efficiently implemented by the licensees.

3.2. Analysis of events at single- and multi-unit sites

Historically, PSAs and risk models have principally been single-unit oriented, where the focus was the specific reactor under analysis, independent from other reactors on site. In recent years, advancements have been made to adapt the classical PSA models to cover multi-units at the same site [29,30]. Studying the nature of events in the database, we observed several instances of multi-unit occurrence, which can be attributed to one or more of the following reasons:

- 1 External events originating outside of the plant boundary that can indiscriminately target one or multiple units on the site.
- 2 Shared components between units of the same site. The components can be part of safety-related systems (emergency power, essential service water), or belong to systems necessary for normal operation of the plant (offsite power, normal service water, condenser water intake, etc.). Failure of these shared components would trigger a simultaneous response from all affected units (e.g. failure of a cross-tied emergency diesel generator).
- 3 Acute or potential failures of non-shared components that, nevertheless, share the same design flaws, inadequate operating procedures, or safety culture deficiencies (e.g. identifying a design problem in the high-pressure injection pump of one unit that is inherent to other pumps of other units on the site).

The database contains events which have occurred in 352 reactors across 143 nuclear sites. Of these sites, 97 (68%) are multi-unit sites that contain 306 reactors. Moreover, around 71% of the 1256 events occurred at multi-unit sites (897 events), with 84% out of these events affecting only one unit on the site, and 14% impacting multiple units simultaneously (Fig. 2). The figure also shows a breakdown of the type of trigger which led to the event, i.e. whether it was an initiating event or a pure system failure; the latter refers to events where no initiating event was observed. The share of the individual initiating events is also displayed. The most common type of trigger affecting multi-units is pure system failures (68%), with the remaining being initiating events (32%). On the other hand, initiating events and pure system failures have equal

Fig. 2. Fraction of affected units at multi-unit sites (left), and analysis of the type of trigger (right): pure system failures (green) and initiating events (yellow). Used abbreviations: LOOP – loss of offsite power; LOCHS – loss of condenser heat sink; IE – initiating event; LOMFW – loss of main feedwater; SBLOCA – small break loss-of-coolant accident; SGTR – steam generator tube rupture.

shares when one unit is affected on the site. What can be immediately observed is that loss of offsite power (LOOP) events dominate the initiating events affecting multiple units, while on the other hand, transients and small break loss-of-coolant accidents (SBLOCA) are more tied to a single unit at the site.

When analyzing the origins of the events (Fig. 3), the majority of initiating events affecting multiple units had external origins (66%), while external events affecting only one unit are at 9%. The share of initiating events originating in the secondary part is also fairly high (27%), which is expected as these include systems shared between multi-units. Finally, we can infer that initiating events originating in the nuclear island are less likely to affect more than one unit (44% for single affected units compared to 7% for multi-units). This breakdown is in line with the aforementioned reasons 1 and 2 (external origins, and shared components respectively).

Shifting our focus to events containing pure system failures on multi-unit sites, only 86 events affected multiple units. Up to 41% of these failures were related to the emergency power and service water systems, which commonly have shared components between units, while the remaining 59% were related to systems usually contained in a single unit. Identifying the contributing factors for failures affecting one or multiple units (Fig. 4.), we can infer that the vast majority of system failures affecting multiple units had causes rooted in design flaws (62%), followed by procedural deficiencies (10.5%). These results confirm our previously identified reasons 2 and 3 for multi-unit occurrences, namely having shared components between the units, or having identical design issues in similar components at multiple units. Regarding single affected units at multi-unit sites, design problems and testing and maintenance (T&M) related errors were the most prominent with 24% each.

The outlined observations can help in the development process of multi-unit PSAs:

- 1 knowing their minor share, this questions the “substantial” need for multi-unit considerations and the assumed interdependencies, and
- 2 the results can orient the focus of multi-unit PSAs towards important initiating events, triggers, origins, systems, and contributors.

3.3. Analysis of common cause failures (CCFs)

Common cause failures (CCFs) are a subset of dependent failures that concurrently result in the unavailability of two or more redundant components/trains of a safety system or function due to the same reason. They can be acute, prompting immediate response from the plant, or potential (latent).² CCFs usually dominate the unreliability/unavailability of safety systems and challenge the effective redundancy. They can be due to one or more of the following reasons:

- Technical factor – redundant trains very often have the same design and operating conditions. This means that, if there is a design problem in one, or some harsh operating conditions, the other trains will have the same problem.
- Human factor – the same maintenance teams in some plants perform their activities on multiple redundant trains, implying that the team is very likely to commit identical errors that threaten the system and lead to a CCF. Additionally, some operator errors have the potential to affect multiple trains.
- Organizational factor – inadequate maintenance and operating procedures usually affect multiple trains. Moreover safety culture deficiencies can affect redundant trains, but even more, a bad safety culture have the potential to affect the whole plant.

In the database, 218 common cause failure occurrences were observed, of which 126 (58%) were acute and 92 (42%) were potential. It should be noted that only failures that were explicitly specified in the events reports to affect multiple redundant trains were counted in this analysis. A breakdown of the detailed contributing factors for these CCF events are presented in Fig. 5. Multiple factors can simultaneously

² Potential CCF: conditional on the occurrence of an event, all redundant trains of a system would have failed, e.g. non-conforming seismic dampers would cause all emergency diesel generators to fail in case of a seismic event.

Fig. 3. Origin of the initiating event affecting: a) multiple units; b) one unit.

Fig. 4. Factors contributing to system failures which have affected multiple units (black), and only one unit (orange), at multi-unit sites.

Fig. 5. Contributing factors to common cause failures (CCF).

contribute to the occurrence of a failure, therefore, in this regard, the factors are not mutually exclusive. Design residuals³ contributed to the majority of CCF events, both acute and potential. Their presence is especially pronounced in potential CCFs with a staggering 70% share of all latent failures. Inadequate procedures were also a very frequent contributor and equally impacting both types of CCFs, while testing and maintenance (T&M) errors⁴ were surprisingly quite common in acute CCFs.

This analysis reaffirms the importance of verification and scrutiny of the design of components, as well as operating procedures. Frequent reviews, inspections, and sharing of operational experience between plants of similar designs – and plants having similar suppliers and manufactures – is necessary in order to limit the future occurrence of CCFs. Additionally, if possible, different teams should be deployed for performing T&M activities on components in redundant trains, to insure some diversification and avoid the occurrence of identical errors due to inadequate knowledge, or failure to follow instructions.

3.4. Quantitative aging signs

Component aging implications have been a growing concern for regulators, plant operators and component manufacturers. Quantitatively, in reliability analysis, one of the most common aging demonstrations is depicted by the bathtub curve of the failure rate of a component [31], which shows high component failure rates during the early periods (infant mortality), a constant rate during the majority of the lifespan, and finally, increased rates at the end of its lifetime (wear-out/aging phase). We anticipate that the cumulative aging of components and their interactions should reveal themselves at a plant level by resulting in signs and more frequent occurrences of events. For this analysis, the following adaptations were performed on the dataset:

- Only events exhibiting contributing factors of technical nature that are likely to be related to aging were taken into account. These include events that contain design-related issues and failures of main and support components. Events that exclusively had failures of human and organizational nature were discarded.
- Events that have simultaneously affected multiple reactors were split into separate events if the reactors were at a different age. The age of the reactor was calculated based on the time of occurrence of the event.

Additionally, a postulated reactor survival rate was calculated based on the experience of 80 permanently shut down reactors that have event entries in the database. Although reactors can be shut down due to multiple reasons (aging, reduced safety, political pressure, financial viability, etc.), this rate calculates the fraction of remaining operating reactors beyond a certain age. Our data shows that around 50% of all reactors are expected to shut down after the age of 30, while less than 10% of reactor are expected to remain operational after the age of 45 (orange curve in Fig. 6).

After calculating the age of a reactor at which an event occurred, the events count per reactor age are plotted in Fig. 6, normalized by the reactor survival rate values. This normalization is necessary in order to get a correct estimate of events counts for every age, by correcting for shut down reactors, and taking into account the survival of reactors beyond a certain age. In other words, the normalization corrects for the fact that a reduced number of operating reactors after a certain age is expected to also yield a reduced number of events. The distribution of

events count vs reactor age exhibits a very interesting behavior (solid dots in Fig. 6). Although not like the classic bathtub curve due to the rather “shy” infant mortality period, the results show an increase in observed events as the reactors age, especially pronounced after the age of 25 years. Reactors past this age appear to be more likely to suffer component failures, or have latent design problems due to technological limits, which also coincides with the usual lifespan of some major components in a nuclear power plant, such as emergency diesel generators, motor-driven pumps, transformers, etc. In conclusion, this analysis shows some traces of aging signs in nuclear power plants after 25 years of operation. Therefore, frequent testing, inspections, and preventive maintenance of components, as well as design reviews and updates must be performed to maintain safe and sustainable operations.

4. Insights and lessons learned from operating experience

Learning from operating experience, efficiently implementing changes, and objectively identifying problems in operating reactors, are very effective when it comes to preventing accidents. After the Chernobyl-4 and the Fukushima Dai-ichi accidents, substantial efforts have been made to increase communication, transparency, and information sharing between countries and regulators in order to continuously increase the level of safety of nuclear power plants. Different countries have different strategies when it comes to identifying problems and implementing changes. An illustrious example is France, where nearly all reactor lines have a similar design and are managed by one operator - Électricité de France (EDF). This is a double-edged sword, as uncovered design issues can plague multiple reactors by default, but on the other hand, addressing them inherently increases the safety of the whole fleet. In the database, at least 15 events can be identified where a problem occurred at one reactor, only for EDF to realize that many reactors in the fleet were affected by the same problem. The majority of these events were related to post-Fukushima safety measures, such as resilience of important safety systems and components to flooding and earthquakes.

Addressing known problems in a timely manner can be invaluable in preventing serious accidents. One shining example from the database is the Fort Calhoun event of January 2010, when external flooding barriers were upgraded after being deemed insufficient. The following year, waves exceeding the height of the initial barriers struck the plant and would have probably caused serious flooding damage if the new barriers had not been installed the previous year. Unfortunately, such nice outcome is not always the rule, as often plants are slow to implement needed changes due to economic considerations, or if the threatening events are deemed unlikely (Fukushima and the TEPCO story serve as an infamous example).

In the following Sections 4.1 to 4.3, we will discuss multiple topics with one common goal – to put emphasis on the importance of learning from operating experience to predict and avoid accidents by implementing changes and exchanging insights with operating peers.

4.1. Precursors to major accidents

In this section, we try to highlight some of the forerunners of the three major nuclear accidents and we discuss how they could have been utilized to foresee and prevent the accidents. We confirm the fact that failing to learn from previous accidents and precursors, and failing to implement respective changes, are the main reasons behind the occurrence of the major accidents [32].

4.1.1. Forerunners and lessons learned from the Three Mile Island 2 accident

One of the most defining moments in nuclear power history is the Three Mile Island (TMI) 2 accident. The accident started at 4 a.m. on March 28, 1979 when loss of main feedwater occurred due to a technical failure. The auxiliary feedwater system was inoperable due to a

³ Design residual CCF example: components (e.g. pumps, valves, etc.) from redundant trains having the same design or manufacturing flaw.

⁴ Testing and maintenance team CCF example: wrong servicing of pumps in redundant trains of a system due to failure in following maintenance instructions.

Fig. 6. Effects of aging in nuclear power plants. Number of events (solid dots) normalized by the postulated reactor survival rate (orange curve).

maintenance error, when all valves were inadvertently left in a closed position. The pilot-operated relief valve (PORV) opened to relieve reactor pressure, but failed to reclose, effectively causing a small break loss-of-coolant accident (SBLOCA). Due to poorly designed instrumentation, the PORV position indicators were ambiguous, causing the operators to believe that there was no leak, and hence, to trip the automatically activated high-pressure injection (HPI) system that provided water makeup to the core, effectively leading to a core dry out and a subsequent meltdown. Fortunately, the next operator shift identified the problem and reactivated the HPI, preventing further core melting. This event had a widespread impact on the industry as a whole, triggering massive design revisions of the plants (primarily control room designs), and discussions about the importance of safety culture, proper training of operating and maintenance personnel, as well as communication between different teams.

In hindsight, there were many instances of similar occurrences (forerunners) around the world which led to less severe outcomes, and lessons stemming from these events were, unfortunately, not well learned and implemented, until the aftermath of the TMI-2 accident. At least three clear forerunners to TMI-2 have been identified, with two of them following a completely identical failure sequence: both Beznau-1, 1974 and Davis Besse, 1977 incidents occurred due to a stuck-open

PORV, and the ambiguous PORV position indicators prevented the operators from correctly identifying the condition. In both of these incidents, the vigilance of the operators prevented further escalation, as they noticed signs of a SBLOCA and acted on their intuition. It is also jarring to see that, one year before the TMI-2 accident, another event occurred at TMI-2, which involved a stuck open PORV and accordingly, a SBLOCA. This event occurred 1 month before the plant was commissioned, presumably during hot functional testing of the reactor. Both Beznau-1 and Davis Besse were designed by American companies, Westinghouse Electric and Babcock & Wilcox (B&W), respectively. The US Nuclear Regulatory Commission (NRC) failed to identify that the ambiguous PORV position indicator is a serious problem that can plague many plants of similar design. More specifically, B&W failed to learn from the Davis Besse incident and failed to implement changes or provide appropriate operator training to other plants sharing their design, including TMI-2.

In addition to the above qualitative discussion, a quantitative analysis have been performed in Fig. 7, showing the cumulative number of events when a SBLOCA was caused by a stuck-open PORV over time in our database. The clear sharp increase in the cumulative number of forerunners before the 1979 TMI-2 accident could be seen as a very good sign of the predictability of the TMI-2 accident (resembling a bubble with an accelerating price in financial markets [33]).

Fig. 7. Cumulative number of events where failed (stuck-open) pilot-operated relief valves (PORVs) resulted in a small break loss-of-coolant accident. Highlighted are the Three Mile Island 2 accident of 1979 (red) and its major forerunners (orange).

Following the TMI-2 accident, we can observe a change of regime, demonstrated with the reduced number of events (plateauing) related to stuck-open PORVs SBLOCA incidents; only 4 events have occurred in the 40 years following the accident (0.1 events/year). This is a decrease by a factor of 10 when compared with the pre TMI-2 period, where 8 events occurred in only 9 years (~ 1 event/year). This encouraging quantitative finding can be attributed to the post TMI corrective design changes, frequent inspections, increased communication levels (with regulators and between countries), and improved training.

4.1.2. Forerunners and lessons learned from the Chernobyl 4 accident

The Chernobyl-4 accident of 1986 was the most severe nuclear accident in the history of civil nuclear power. The RBMK light water graphite reactor (LWGR) design was inherently flawed in multiple aspects: a positive void coefficient; a positive scram effect where control rods had a graphite tip, which, although useful for normal operation for smoothing out the neutron flux, inadvertently aided in neutron moderation thus increasing reactivity; the core was massive, with uneven flux distribution and physical conditions that were difficult to model; instrumentation for reactivity measurement were unavailable for values below 10% of rated power; there was no containment to prevent radiation release in case of an accident. Additionally, Soviet plants at the time had a seriously deficient safety culture, organizational procedures, and personnel training: lack of transparency and censoring of outspoken critics of the design led to a false belief of “imperviousness” of the reactors; unavailability of decent modeling software led to a poor understanding of phenomena such as fuel burnup; the staff rarely questioned supervisors due to fear of consequences.

The accident started around 1 a.m. with a turbo-generator rundown test, one of the important emergency operating modes. After the start of the test, the reactor power suddenly dropped, prompting the operators to raise completely some of the control rods in order to increase the power. Unbeknownst to them, the low power operation and high fuel burnup led to a build-up of the neutron poison Xenon-135, which was the reason why the power was unresponsive to the lifting of control rods. After the power increased, the test commenced at 1:22 a.m. and rundown was performed on 4 main feedwater pumps, which reduced the coolant flow. This resulted in a sudden power increase (positive void coefficient), prompting the operators to drop the control rods (having a positive scram deficit), which led to an instantaneous power excursion. The uncontrollable reactivity and the drastically increasing temperature and pressure ultimately led to core meltdown and a large steam explosion, which destroyed the reactor core and the building, followed by graphite burning that contributed to a wide-spread release of radioactive isotopes.

Due to the lack of transparency, the exact number of Chernobyl-4 forerunners is unknown. In our database, we have 2 events that appear to meet the criteria, although information about them is also limited. These events are Leningrad-1, 1975 and Beloyarsk-2, 1977, both of which suffered a melting of fuel channels (one channel at Leningrad, and approximately half of the core at Beloyarsk), due to a loss of coolant, presumably aggravated by the positive void coefficient issue of the RBMK design. One of the hypothesized reasons for the event at Beloyarsk is testing of a new cooling water purification system, although no official explanation was given by the Soviet Union. Needless to say, the design and organizational issues brought up before were well known by the regulatory body, however, they failed to act in time to avoid the occurrence of more serious accidents. This further stresses the importance of transparent reporting, communication between plant operators and regulators, and generally learning from operating experience. The effects of the Chernobyl accident were far-reaching in the nuclear world, spawning a period of heightened awareness, panic reporting, increase of transparency, and major organizational and design changes for RBMKs.

4.1.3. Forerunners and lessons learned from the Fukushima Dai-ichi accident

The Fukushima Dai-ichi nuclear accident of March 11, 2011 was a grim reminder of the devastating potential of beyond design-basis accidents. The accident occurred as an aftermath of a tsunami triggered by the magnitude 9 Tohoku Earthquake. Before the event, units 1, 2 and 3 were in normal power operation, unit 4 was refueling with unloaded fuel, and units 5 and 6 were under cold shutdown. The earthquake caused a loss of offsite power to all units, and reactors 1, 2 and 3 were tripped. Emergency diesel generators (EDGs) started and provided emergency power to all units. Units 2 and 3 were cooled using the reactor core isolation cooling (RCIC) system, while the unit 1 used the isolation condenser, which was cyclically operated to cool down the reactor. The isolation condenser tripped when the tsunami occurred, and stayed isolated due to lack of any power. The tsunami hit shortly afterwards, with waves exceeding the protection of the plant, completely flooding the EDGs and service water systems and resulting in a station blackout. The RCIC continued operating in units 2 and 3 but was tripped eventually due to loss of DC power control, and units 1, 2 and 3 started overheating. Units 5 and 6 were newer, having higher flood barriers, therefore, one EDG in unit 6 managed to survive and provided power to the cooling systems of both units through a cross-tie. The plant operator Tokyo Electric Power Company Holdings (TEPCO) failed to bring external EDGs to the site on time, resulting in a core melt in units 1, 2 and 3. Hydrogen build up resulted in an explosion and breach of containment, and radioactive isotopes were released in the surrounding area. Thousands of people were ultimately displaced, although no casualties have been reported as a direct consequence of the accident thus far. Three other plants in Japan were also affected by the Tohoku Earthquake and the tsunami (see Fig. 8), meaning that a larger catastrophe could have occurred, but fortunately, the EDGs in these plants were not completely flooded.

Despite the beyond design-basis nature of the trigger of the Fukushima Dai-ichi accident, steps could have been taken to avoid or at least mitigate the effects of such a devastating natural phenomenon. In our database, we have mapped out several events that can be considered forerunners to the Fukushima accident, where lessons were not well learned and implemented. Two events are considered to be flooding-related forerunners, while we would also like to put emphasis on two additional events that can be seen as forerunners in a station blackout sense (Fig. 8). The flooding events are the Blayais incident of 1999, and Madras-2, 2004. In both events, the flooding protection of the plants was insufficient, leading to multi-component failures and threatening the cooling capabilities of the reactors. In Blayais, the offsite power supply was lost but, fortunately, the EDGs were unaffected by the flooding. Madras suffered an identical scenario as Fukushima, with a 9.1 earthquake triggering a tsunami that overwhelmed the protection of the plant. Fortunately, only the condenser and seawater pumps were lost, while the EDGs and service water system remained operational. These serious precursor events should have triggered an industry-wide debate regarding the integrity of coastline nuclear plants and the threats of beyond design tsunamis, as well as the proper protection of vital equipment. Additionally, knowing that the newer units on the Fukushima site – namely units 5 and 6 that managed to avoid core damage – had higher flood barriers, raises the question if the plant operator and regulators were aware of the potential flooding vulnerability of older units, knowing that the vital equipment were not properly protected in the event of any flooding exceeding the design basis.

Furthermore, the station blackout events of Kola 1–2, 1993 and Maanshan 1–2, 2001 are emphasized because both of these plants avoided core damage thanks to attaching external (swing) power sources (EDGs, reversible motor-generators) after a station blackout occurrence. These events demonstrated the importance of having available, and easy to deploy, power sources in the near vicinity of the plant, but still far enough to not be affected by the same external initiators. Fukushima Dai-ichi, unfortunately, did not have such power sources

Fig. 8. Forerunners to the Fukushima Dai-ichi accident. Events shown in the figure: Tohoku Earthquake events of March 11, 2011 are Fukushima Dai-ichi 1, 2, 3, 4, 5, 6; Fukushima Dai-ichi 1, 2, 3, 4; Tokai 2; and two events at Onagawa-2. Station blackout events are the Kola 1, 2, 1993; and Maanshan 1, 2, 2001.

available, and the lack of communication/transparency between TEPCO and the government resulted in a failure to bring additional EDGs to the site on time.

The effects of the Fukushima Dai-ichi accident and its lessons had a major impact on the nuclear industry. Multiple design changes related to natural hazards were implemented around the world, such as the stress tests within Europe and the external power sources retrofits. Work has also been done in improving transparency and communication between companies and regulatory bodies, as well as with the general public.

4.2. Primary water stress corrosion cracking (PWSCC) events

Component failures rooted in the design, choice of materials, and manufacturing process, always have the potential to affect multiple plants. This was the case with the clustering of primary water stress corrosion cracking (PWSCC) events identified in the 80's and early 2000's in the United States. The three major clusters that can be observed are: PWR main feedwater events, BWR recirculation piping events, and PWR reactor pressure vessel (RPV) leakage events (Fig. 9).

The root cause for all these PWSCC clusters was the usage of material that was susceptible to stress corrosion (due to lack of knowledge), although their implications varied significantly.

The PWR main feedwater events cluster mostly included acute occurrences, defined with rupture of pipes in the heat extraction and reheater lines causing immediate plant shutdowns. Only the last event (Trojan, 1987) was prevented before it occurred, as the maintenance team identified signs of accelerated corrosion on the main feedwater pipes and immediately took action. This is a good indication of learning from experience, as Trojan already suffered a similar problem in 1982, so the maintenance team was aware of the problem and worked on preventing its reoccurrence. It is unknown if these phenomena extended beyond the three affected plants shown in Fig. 9 (Trojan, Oconee, Zion).

The BWR recirculation piping degradation events were first identified in Monticello and Nine Mile Point plants in 1982. In Nine Mile Point-1, leaks were observed from the piping, so the plant operators immediately shut down the reactor and opted to replace the pipes. This prompted a series of investigations across multiple BWR plants where the same material was used, leading to the later discoveries presented in

Fig. 9. Clustering of events caused by leakage due to PWSCC. PWR main feedwater events: Oconee-2, 1982; Zion-1, 1982; Trojan, 1982, 1987. BWR recirculation piping events: Nine Mile Point-1, 1982; Monticello, 1982; Pilgrim, 1983; Browns Ferry-3, 1983; Peach Bottom-2, 1985. PWR RPV events: Oconee-1, 2000; ANO-1, 2001; Crystal River-3, 2001; North Anna-2, 2001; Oconee 2, 3, 2001; Palisades, 2001; Surry-1, 2001; Three Mile Island-1, 2001; Davis Besse, 2002; ST. Lucie-2, 2003.

the figure. In total, 5 plants were affected and costly repairs were performed to remedy the problem. Knowing the potential severity of a recirculation piping rupture (dangers of LOCA), this cluster of events serves as a good example of how transparency and good communication between plant operators and the regulatory body can trigger reviews, and aid in preventing serious incidents.

By far, the most significant of the clusters is the RPV leakage events in the early 2000's. Identified in 11 reactors, this "pandemic" of events involved leakage from the control rod drive mechanism nozzles due to PWSCC. PWRs use borated water for neutron flux control, which can quickly corrode susceptible material, as was the case in these events. This material susceptibility problem could have had disastrous consequences as the structural integrity of the reactor pressure vessel was threatened. The most serious of the events, the Davis Besse "near-miss" of 2002, had immense corrosion damage on the RPV head due to borated water leakage from a cracked control rod drive mechanism – only 9.5 mm of steel cladding were left, holding the high pressure (17 MPa) coolant in the vessel. If the vessel had been breached, a large break loss-of-coolant accident (LBLOCA) would have occurred, debris could have clogged the emergency core cooling system, and the ability of the plant to combat this situation would have been put into question. Fortunately, this was avoided as the problem was identified during the scheduled inspections, triggered after several precursory events had occurred in early 2001 (Fig. 9). The good communication and transparent approach between regulators and operators of the plants prevented any accident from occurring.

4.3. Interesting events insights based on operating experience

The information presented so far has been largely based on observations of known accidents, problems, systems, and phenomena. However, the database contains some important events that are interesting, strange, unanticipated, or generally difficult to model using the standard PSA models. These observations and events have important takeaways that will be emphasized in this section.

4.3.1. Propagating events

The importance of redundancy of individual trains was already discussed and how steps must be taken to prevent failures to propagate within the system (dependent and common cause failures). However, inter-system and inter-unit failure propagations can also occur, and the consequences can be quite serious. This behavior is difficult to model in standard PSAs and, based on our experience, usually occurs due to latent design errors and component failures. Two events will be presented that have inter-unit propagation, and one intra-unit event having unexpected interaction between different systems.

(Smolensk 2, July 1990)

During testing the automatic activation of a unit 1 emergency power source, the maintenance team inadvertently damaged the cable insulation, causing a short circuit and a fire. The fire damaged unit 2 cables that were laid nearby, causing a loss of two emergency power buses, loss of power to two reactor coolant pumps and a reactor scram. The placement of cables in close proximity to each other for both units was considered to be a serious design flaw.

(Calvert Cliffs-2, February 2010)

One reactor coolant pump (RCP) at unit 1 tripped due to a ground fault. The breakers failed to isolate the ground fault close to the source, which then propagated and tripped the station auxiliary transformer (SAT). All RCPs and condenser recirculation pumps at unit 2 lost power, resulting in an automatic reactor trip from 99.5% power. The loss of the SAT additionally resulted in a loss of the power to one of the safety buses, prompting one emergency diesel generator (EDG) to start. However, it tripped shortly after due to a low lubrication oil pressure. The reactor was ultimately stabilized.

4.3.2. Challenge to the Swiss cheese model

The Swiss cheese model has been always used to highlight the importance of several layers of defense, and the unlikelihood of accidents in complex and reliable technologies. Although to a large extent, this is a very representative model, nevertheless, accidents and events have shown that anomalies can align, even independently, forming a challenge to the multi-layers model. Here, we give some prominent examples of interesting events from our database, where multiple failures happened independently, questioning the unlikelihood of these multiple "unlikely occurrences". Additionally, this brings questions regarding trusting the individual probabilities and numbers, independence assumptions, validity of reductionism, the role of feedback loops, and others.

(Brunswick, April 1975)

At 10-percent power, a safety relief valve (SRV) was stuck open, rendering the reactor core isolation cooling (RCIC) system inoperable. Operators failed to scram the reactor according to the emergency operating procedures. The high-pressure coolant injection (HPCI) system failed to run as a result of high torus level. Loop B of the residual heat removal (RHR) system failed due to a failure in its respective service water supply. The reactor scrammed automatically after a manual closure of the main steam isolation valve (MSIV).

(Davis Besse, September 1987)

While the plant was in full power operation, a failure in the main feedwater flow transmitter caused an increased flow to the steam generators, causing reactor overcooling and reactivity increase. The operators wanted to insert some control rods to reduce and control the power. However, they inadvertently inserted the axial power shaping rods, causing a power spike; the reactor tripped on the high neutron flux signals. During the event, other unrelated anomalies also occurred: a breaker failure caused a loss of power to one safety bus, a service water pump failed to automatically start due to a maintenance error, an atmospheric dump valve and a turbine bypass valve failed open causing excessive cooling, finally, the RHR system was slightly affected due to a void in the discharge piping owed to a nitrogen leakage.

4.3.3. Plant level common cause failures

Systems reliability is known to be threatened by common-cause failures, where CCFs dominate the unreliability quantification of a redundant system. Nevertheless, thanks to experience, CCFs at components level have been well understood and taken into account in reliability analysis, and systems are designed in a way to minimize the interactions between components and common-cause type of failures through the different concepts of diversification. However, some events in the database surprisingly show signs of a more serious type of CCFs, i. e. at a higher level, namely, at system/plant-level. Plant-level CCFs describe failures and interactions between different independent systems, which are not typically taken into account in PSA modeling. In this section, we give some examples of events showing plant-level CCF potential, which could eventually dominate the core damage sequences.

(COOK, October 1999)

An evaluation of the high-energy line break (HELB) program identified important equipment that were not qualified for the harsh environment resulting from a HELB, or would have been damaged by jet impingement from the break. Potentially affected systems include: auxiliary feedwater (AFW), safety and non-safety-related 600 VAC and lower voltage switchgears, emergency diesel generators (EDGs), component cooling water (CCW), auxiliary building ventilation equipment, cables and conduits inside the containment, and other equipment located near 1- the pressurizer surge line, 2- the chemical volume control system (CVCS) letdown, 3- the steam generator blowdown piping, and 4- HELB doors. The cause of this plant-level CCF is the fact that dynamic loads have not been adequately taken into account; the event triggered several plant modifications and HELB procedural changes.

(Cooper, June 2007)

Multiple systems potentially affected, specifically, the residual heat

removal system and the high-pressure coolant injection, due to inadequate post-fire operating procedures.

4.3.4. Unusual errors during testing and maintenance activities

Understanding human related failures is difficult due to its complex nature, and it is the basis of human reliability analysis. In this paper and in [26], the contributions of human related failures were discussed, showing that they very often have a significant, and sometimes a dominating, share of the contributing factors, especially failures related to T&M activities. This section presents some completely unusual examples of maintenance errors, where the team could not foresee the consequences of their actions, or the lack of focus or information have led to bizarre occurrences.

(Davis Besse, April 1980)

During cold shutdown, core cooling was provided by residual heat removal (RHR) pump 2, with pump 1 being unavailable due to maintenance. Mechanical vibrations due to maintenance work in the switchgear room suddenly caused a breaker ground relay to actuate, causing two essential buses to trip. The loss of the buses caused problems in the RHR pump 2, as the logic circuit transferred the suction of the pump to the sump, which was empty. This resulted in air being drawn into the suction of the pump, which tripped and resulted in the loss of reactor cooling. The pump was recovered after 2.5 h and reactor cooling was re-established.

(Brokdorf, July 1995)

While the plant was shut down and the core was unloaded to the spent fuel pool, the T&M team erroneously started activities on emergency diesel generators (EDGs) 4 and 8 instead of the designated 1 and 5, making them inoperable. This condition did not cause acute distress to any safety system, but a loss of offsite power event would have led to loss of spent fuel cooling. However, the long grace period for manually starting the EDGs and the additional measures in place would have been sufficient to quickly re-establish the spent fuel cooling.

(Pickering-5, July 1995)

In a bizarre turn of events, two technicians carried out maintenance work on the fast shutdown system of the wrong reactor (reactor 5 instead of reactor 6), disabling the system for the reactor that was operating at full power at the time.

(Monticello, December 2014)

During testing of one emergency diesel generator (EDG), a non-licensed operator erroneously adjusted the settings of the governor belonging to the other EDG. This effectively caused both EDGs to be declared inoperable.

5. Conclusion

With more than 1250 insightful events from commercial nuclear power plants, the ETHZ curated nuclear events database is the largest open-access database in the world. In addition to the constructive discussions offered in [26], the present paper has presented important analyses and industry insights with the goal of raising awareness to several key aspects threatening the safety and reliability of nuclear power plants.

Important takeaways, which can be highlighted from the discussions, include:

- Major accidents always trigger a wave of “reactive” reporting, mostly due to increased alertness, improved transparency, uncovering latent design errors, and heightened public pressure. This trend appears to last 5 to 6 years after which the implemented design changes inevitably lead to safer operation and lower number of events. Additionally, changes in regulatory or corporate management, which usually takes place in cycles of 5–6 years, can result in lower alertness and weakened transparency culture, and lead to deteriorating reporting habits. A good approach that can be borrowed from the military is the war drill practice, i.e. in order to keep the level of

vigilance high requires to constantly challenge the peacefulness of operation, and to have accident mockups and simulations, which can keep the plant operators and personnel alert at all times.

- The leading causes for multi-unit events are external triggers and design issues. These factors must be taken into account the most, when adapting PSAs to cover multi-unit events.
- Common-cause failures (CCFs) occur fairly frequently and can completely erode the safety of the plant. Frequent review of component design and operating procedures, employing different teams for T&M activities on redundant trains, and sharing operational experience between plants of similar designs, are some of the steps that should be taken in order to limit future occurrences of CCFs.
- Quantitative signs of aging can be observed for plants after the age of 25. Therefore, frequent testing, inspections, and design reviews and updates should be performed by the licensees.
- Learning from previous accidents and precursors is essential in preventing the occurrence of new accidents. This calls for a larger recording, reliance, and sharing of operational data to support learning from experience and avoid reoccurrence of accidents and events.
- Important lessons can be learned from events, as qualitative insights can be very beneficial. Raising awareness to these issues can potentially bring light to similar vulnerabilities in currently operating plants.

The ETHZ nuclear events database can be tremendously beneficial to many interested parties, including researchers, industry partners and the general public. We trust that, with increased regulations for operational monitoring, events recording, increased automation and sensor technology levels, aided by big data and machine learning capacities, we are reaching a level of unprecedented abilities for learning and predicting. The structure and utilization of the database can serve as best practice for other areas in industry. The database, along with the access tool and user manual, are publicly available at: <http://er-nucleardb.ethz.ch/>.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] Kröger W, Zio E. *Vulnerable systems*. Springer Publishing Company; 2011. Incorporated.
- [2] Sornette D, Kröger W, Wheatley S. *New ways and needs for exploiting nuclear energy*. Springer; 2018.
- [3] Mosleh A. PRA: a perspective on strengths, current limitations, and possible improvements. *Nuclear Eng Technol* 2014;46(1):1–10.
- [4] Kröger W, Sornette D. *Reflections on Limitations of Current PSA-Methodology*. International Topical Meeting on Probabilistic Safety Assessment and Analysis (PSA). 2013.
- [5] Zhang X, Mahadevan S. Bayesian network modeling of accident investigation reports for aviation safety assessment. *Reliab Eng Syst Saf* 2021;209:107371.
- [6] Birkhofer A, Köberlein K. *Data situation and the quality of risk assessment, in reliability data collection and use in risk and availability assessment*. Springer; 1986. p. 3–15.
- [7] Kröger W, Sornette D, Ayoub A. *Towards Safer and More Sustainable Ways for Exploiting Nuclear Power*. *World J Nuclear Sci Technol* 2020;10(03):91.
- [8] Duffey RB. *The risk of extended power loss and the probability of emergency restoration for severe events and nuclear accidents*. *J Nuclear Eng Radiat Sci* 2019; 5(3).
- [9] Mach P. *Life management and operational experience feedback-tools to enhance safety and reliability of the NPP*. In: *Proceedings Of The 2nd International Symposium On Safety And Reliability Systems Of Pwrs And Vvrs*. Czech Republic; 1997.
- [10] Rasmussen J. *Risk management in a dynamic society: a modelling problem*. *Saf Sci* 1997;27(2):183–213.
- [11] Khakzad N, Khan F, Paltrinieri N. *On the application of near accident data to risk analysis of major accidents*. *Reliab Eng Syst Saf* 2014;126:116–25.

- [12] IAEA, Nuclear power plant operating experience. 2018, International Atomic Energy Agency: Vienna.
- [13] Moura R, Beer Michael, Patelli Edoardo, Lewis John, Knoll Franz. Learning from accidents: interactions between human factors, technology and organisations as a central element to validate risk studies. *Saf Sci* 2017;99:196–214.
- [14] Wheatley S, Sovacool B, Sornette D. Of disasters and dragon kings: a statistical analysis of nuclear power incidents and accidents. *Risk Anal* 2017;37(1):99–115.
- [15] IAEA, IAEA releases 2019 data on nuclear power plants operating experience. 2020.
- [16] Ayoub A, Stankovski S, Kröger W, Sornette D. Status of the ETHZ Curated Nuclear Events Database. 30th European Safety and Reliability Conference and 15th Probabilistic Safety Assessment and Management Conference (ESREL2020 PSAM15), Venice, Italy. 2020.
- [17] Ayoub A, Stankovski Andrej, Wheatley Spencer, Kröger Wolfgang, Sornette Didier. ETHZ Curated Nuclear Events Database. 2020. retrieved from <http://er-nucleardb.ethz.ch/>.
- [18] Ayoub A, Kröger W, Nusbaumer O, Sornette D. Simplified/Harmonized PSA: a Generic Modeling Framework. International Topical Meeting on Probabilistic Safety Assessment and Analysis. Charleston, SC, USA: American Nuclear Society (ANS); 2019.
- [19] Ayoub A, Kröger W, Sornette D. Generic Probabilistic Safety Assessment Models for International Precursor Analysis Applications. In: International Youth Nuclear Congress (IYNC); 2020.
- [20] U.S.NRC., Licensee Event Report Search (LERSearch). 2020: retrieved from <https://lersearch.inl.gov/Entry.aspx>.
- [21] U.S.NRC., Accident Sequence Precursor (ASP) Program. 2020: retrieved from <https://www.nrc.gov/about-nrc/regulatory/research/asp.html>.
- [22] Johnson JW, Rasmuson DM. The US NRC's accident sequence precursor program: an overview and development of a Bayesian approach to estimate core damage frequency using precursor information. *Reliab Eng Syst Saf* 1996;53(2):205–16.
- [23] WANO, Performance Analysis Program. 2020: retrieved from <https://www.wano.info/services/performance-analysis>.
- [24] IAEA, IRS Guidelines – Joint IAEA/NEA International Reporting System for Operating Experience. 2010: Vienna.
- [25] JRC E. Clearinghouse on Operating Experience Feedback 2020. retrieved from https://clearinghouse-oef.jrc.ec.europa.eu/search/oef_records.
- [26] Ayoub A, Stankovski A, Kröger W, Sornette D. The ETH Zurich curated nuclear events database: Layout, event classification, and analysis of contributing factors. *Reliab Eng Syst Saf* 2021;213. <https://doi.org/10.1016/j.res.2021.107781>.
- [27] Chernov D, Ayoub Ali, Sansavini Giovanni, Sornette Didier. *Averting Disaster Before It Strikes: how to ensure your subordinates warn you while there is still time to act*. Springer; 2022. Handbook to be published.
- [28] Chernov D, Sornette Didier, Sansavini Giovanni, Ayoub Ali. *Disastrous Codes of Silence: how poor communication within organizations causes major catastrophes*. Springer; 2021. Book to be published.
- [29] Lim H-G, Kim Dong-San, Han Sang Hoon, Yang Joon Eon. Development of logical structure for multi-unit probabilistic safety assessment. *Nuclear Eng Technol* 2018; 50(8):1210–6.
- [30] Modarres M, Zhou T, Massoud M. Advances in multi-unit nuclear power plant probabilistic risk assessment. *Reliab Eng Syst Saf* 2017;157:87–100.
- [31] Birolini A. *Reliability engineering: theory and practice*. Springer Science & Business Media; 2013.
- [32] Dechy N, Rousseau J-M, Jeffroy F. *Learning lessons from accidents with a human and organisational factors perspective: deficiencies and failures of operating experience feedback systems. nuclear safety: new challenges*. Paris: gained experience and public expectations; 2011. 7.
- [33] Sornette D. *Why stock markets crash: critical events in complex financial systems*, 49. Princeton University Press; 2017.