1/12/2024

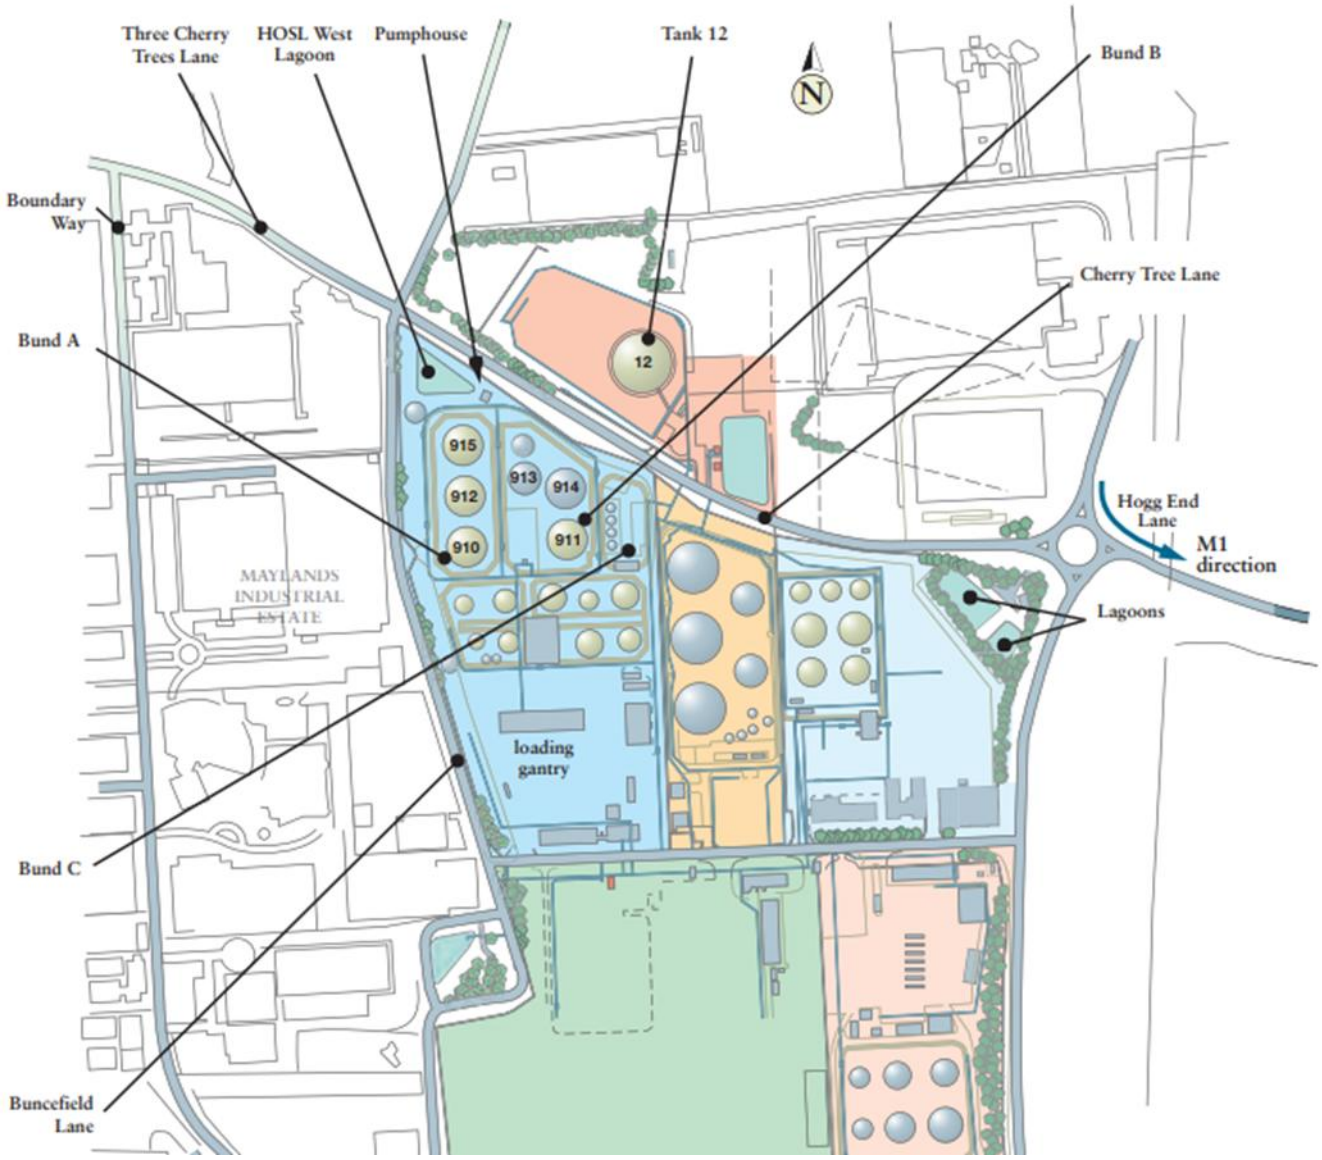

From - *The Buncefield Incident 11 December 2005: The Final Report of the Major Incident Investigation Board*. Health and Safety Executive.

CAMBRENSIS

The Buncefield Disaster: A Case Study in Safety Failures and Systemic Oversight

The Paradox of PFD's | David Slater

THE BUNCEFIELD DISASTER: A CASE STUDY IN SAFETY FAILURES AND SYSTEMIC OVERSIGHT

David Slater, Cardiff University, dslater@cambrensis.org

THE INCIDENT.

On December 11, 2005, the Hertfordshire Oil Storage Terminal near Hemel Hempstead, England, became the site of one of Europe's most devastating industrial accidents. (1) Known as the Buncefield disaster, this event resulted from a cascade of technical malfunctions, human oversights, and organizational failures. The explosion and ensuing fires caused extensive damage, widespread environmental contamination, and a review of safety practices in the petrochemical storage industry. The Buncefield incident has since become a case study in industrial safety, illustrating how the failure of preventive and mitigative barriers can lead to catastrophic consequences when exacerbated by systemic issues and operational gaps.

The disaster began on December 10, 2005, when unleaded petrol started filling Tank 912 at a rate of 550 cubic meters per hour. This process continued into the early hours of December 11, despite several critical issues. At approximately 03:00, the tank's automatic level gauge froze, showing a constant reading and providing no indication that the tank was nearing capacity. By 05:20, the tank began to overflow as petrol spilled out through the roof vent into the surrounding bund, a containment area designed to hold spills. The bund, however, was not capable of handling the volume of petrol that continued to pour out unchecked. CCTV footage later showed a dense petrol-air vapour cloud forming and spreading beyond the bund and into adjacent areas. The situation worsened further when the inflow rate increased to approximately 890 cubic meters per hour.

Despite alarms and the visibly growing hazard, the pumping continued until 06:01, when the vapour cloud ignited. The explosion measured 2.4 on the Richter scale and was heard as far as 125 miles away. The blast engulfed 20 large storage tanks in flames, leading to a fire that burned for several days and produced a smoke plume visible from space. While the explosion caused extensive damage, the timing of the incident—early on a Sunday morning—meant that no operators were in the immediate vicinity, and the nearby Maylands Industrial Estate was unoccupied. This fortuitous timing likely prevented loss of life, though 43 people sustained injuries, including two seriously.

Although the Buncefield oil storage terminal was not entirely remotely operated, its operations relied heavily on automated systems for monitoring and control. Automated systems were responsible for overseeing tank levels, flow rates, and triggering alarms in case of anomalies. The level gauge on Tank 912, a critical part of the automated system, failed early in the process, freezing at a constant reading. This malfunction gave operators no indication that the tank was nearing capacity, and no manual checks were performed to verify the automated readings. The independent high-level safety switch, which was designed to halt inflow when the tank reached capacity, was also disabled or inoperative, further compounding the problem.

Operators on-site relied heavily on these faulty automated systems without conducting manual verifications. This reliance was symptomatic of a broader issue: a systemic dependence on automation without adequate redundancy or human oversight. Investigators later noted that the timing of the incident—early on a Sunday morning—likely coincided with reduced staffing levels, limiting opportunities for human intervention. Even if alarms were triggered, there was no evidence that they prompted action. This failure suggests either the alarms were not designed to capture sufficient attention, or operators lacked the training to respond effectively to such warnings.

The sequence of failures reflects a breakdown of safety barriers at multiple levels. The level gauge on Tank 912, which had a history of malfunctions, was not repaired or replaced despite clear evidence of unreliability. Maintenance records showed that the high-level safety switch had been inoperative since August 2005, yet no corrective action was taken. Alarms, designed to warn of high levels or potential overflows, either failed to activate or were ignored due to operator inattention or inadequate training in alarm management protocols. The bund, a secondary containment measure intended to prevent spills from spreading, was poorly maintained and incapable of containing the volume of petrol released. It allowed the vapour cloud to escape into the surrounding environment, where it posed an ignition risk.

Operator oversight was limited, and manual checks were not conducted despite the known issues with the automated systems. Investigations found no evidence that operators were asleep or absent during the event, but their reliance on faulty systems and lack of intervention highlighted significant gaps in training and operational procedures. Fatigue may also have played a role, as the incident occurred during the early morning hours, a time when cognitive performance is naturally reduced.

A Post-Facto Analysis

A Bow Tie analysis of the Buncefield disaster provides a structured framework to understand the threats, barriers, and failure modes associated with the event. At the centre of the analysis is the release of a flammable vapour cloud due to the overfilling of Tank 912. Preventive barriers, including the level gauge, high-level safety switch, alarms, and operator monitoring, all failed. Mitigative barriers, such as the bund, ignition source control, and firefighting systems, were either overwhelmed or insufficiently designed to manage the scale of the disaster. The vapour cloud's ignition, likely caused by a static discharge or other unprotected spark, initiated the explosion, which in turn overwhelmed local firefighting resources and caused extensive damage to surrounding properties.

Examples of such Bow tie analyses are referenced below (2,3,), and the CCPS Bow Tie (Figure 1) is perhaps the clearest)

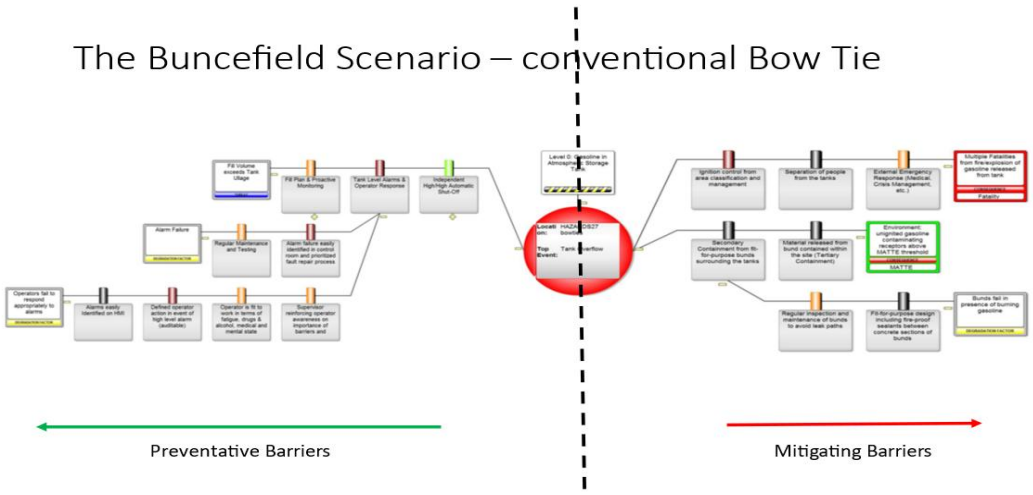

Fig. 1 – The CCPS Bow Tie for Buncefield

How safe was it?

The company would probably have felt it was safe enough and complied with the relevant standards. They would have assumed that the "barriers" in place were sufficient to guarantee safe operation *As imagined". A standard system safety approach could have been used to assign a probability of Failure on demand (PFD), to each of the barriers and use the Layers of Protection Analysis (LOPA) method of predicting their combined effectiveness as a System Integrity level (SIL). (4).

These barriers were notionally, in sequence, the level gauge, the High -Level switch, the Alarms, the monitoring operator, the containment systems, the ignition source control, the firefighting systems, and probably the spacing and location of the tanks to prevent fire spread and external damage.

Applying some ballpark estimates for these numbers we get –

> **SIL** (as imagined) = (Gauge)10-3 x (Switch)10-3 x (Alarm)10-2 x (Operator)10-1x (Containment)10-1 x (Ignition)10-1 x (Firefighting)10-1 x (Spacing)10-2 = 10-14

Safe enough? (once in the age of the universe?). But in fact, these barriers were not at the standard level of reliability envisaged.: -

The Preventive Barriers

- **Level Gauge:** Failed due to technical malfunction and maintenance neglect. The level gauge on Tank 912, a critical tool for monitoring fuel levels, had a history of malfunctions. Despite this, operators relied on it without cross verifying its readings.
- **High-Level Safety Switch:** Similarly, the independent high-level safety switch, designed to stop fuel inflow when the tank reached maximum capacity, was disabled or inoperative, rendering it ineffective. Maintenance logs revealed that it had been unreliable since its last servicing in August 2005, yet no corrective action was taken.
- **Alarms:** Alarms tied to high fuel levels or vapour detection either failed to activate or were not adequately monitored. Operators missed opportunities to address the escalating

hazard, possibly due to reliance on malfunctioning systems and inadequate training in alarm escalation protocols.

- **Operator Monitoring:** Lacked manual checks and redundancy in monitoring systems. Operators did not detect the overflow in time, largely due to over-reliance on automated systems and insufficient manual checks. The timing of the incident in the early morning hours raises questions about vigilance during night shifts and the potential role of fatigue. Investigations found no evidence of operators being asleep or absent, but systemic issues such as inadequate training and alarm management likely contributed to the oversight.

The Mitigative Barriers

- **Containment Systems:** The bund was inadequate for the scale of the spill. The bund around Tank 912 was a secondary containment measure, intended to prevent fuel from spreading. However, it was not designed to manage the sheer volume of the overflow. Poor maintenance allowed the vapour cloud to escape the containment area, exacerbating the risk of ignition. The sealing materials would probably have melted in the fire.
- **Ignition Source Control:** An offsite source ignited the vapour cloud, highlighting gaps in hazardous area management.
- **Firefighting Systems:** Overwhelmed by the intensity and scale of the fire.
- **Separation distances:** The tank spacing was not as recommended and a dense, passive unconfined, vapour cloud explosion (UCVE), of cold petroleum had not been documented before. (5)

SIL (as done) = 1x1 x 1 x 0.5 x 1 x 1 x 0.8 x 1 = 0.4 (40% chance of an incident?)

Could they have predicted this? This raises some legitimate questions about the use of this type of solution to assuring the safety of complex systems, (5), to stack layer upon layer of safety critical subsystems to add more and more confidence in their safer operation. At the Sizewell B inquiry, it was shown that adding an extra software layer actually increased the probability of failure. (6)

THE PARADOX OF PFD

Appendix A attempts to show the problems with calculating realistic levels of system integrity using the LOPA approach. The method treats the barriers as separate independent hermetically sealed black boxes, insulated from whatever is going on in the rest of the system. This may be convenient but is it correct? But not only incorrect but misleading and giving a false sense of security. One way perhaps to address this is to use a Bayesian approach to combining the probabilities of performance of individual barriers dependent on the evidence of the reliability of the other safety critical barriers. If they are totally dependent, as in common power supplies, then they obviously have a common failure mode. But what if it is more subtle, as in the quality of the equipment maintenance? Here using Bayesian theory to combine probabilities would be more accurate. With a common dependence, the PFD of a succeeding Barrier will be 1, (Appendix A).

Using the Bayesian adjustment then gives a much more realistic assessment of the probability of such an incident as Buncefield experienced.

BUT SAFETY IS A TOTAL SYSTEM ATTRIBUTE.

Applying a Bayesian approach to separate component Barriers is an improvement, but if the problem is the interdependence and interaction of the barriers with the rest of the system, perhaps we should employ a bona fide system approach? With that in mind the Bow Tie was used to produce a FRAM model using the Functional Resonance Analysis method (Figure 2) (7,8).

A Functional Bow Tie of the Buncefield Incident

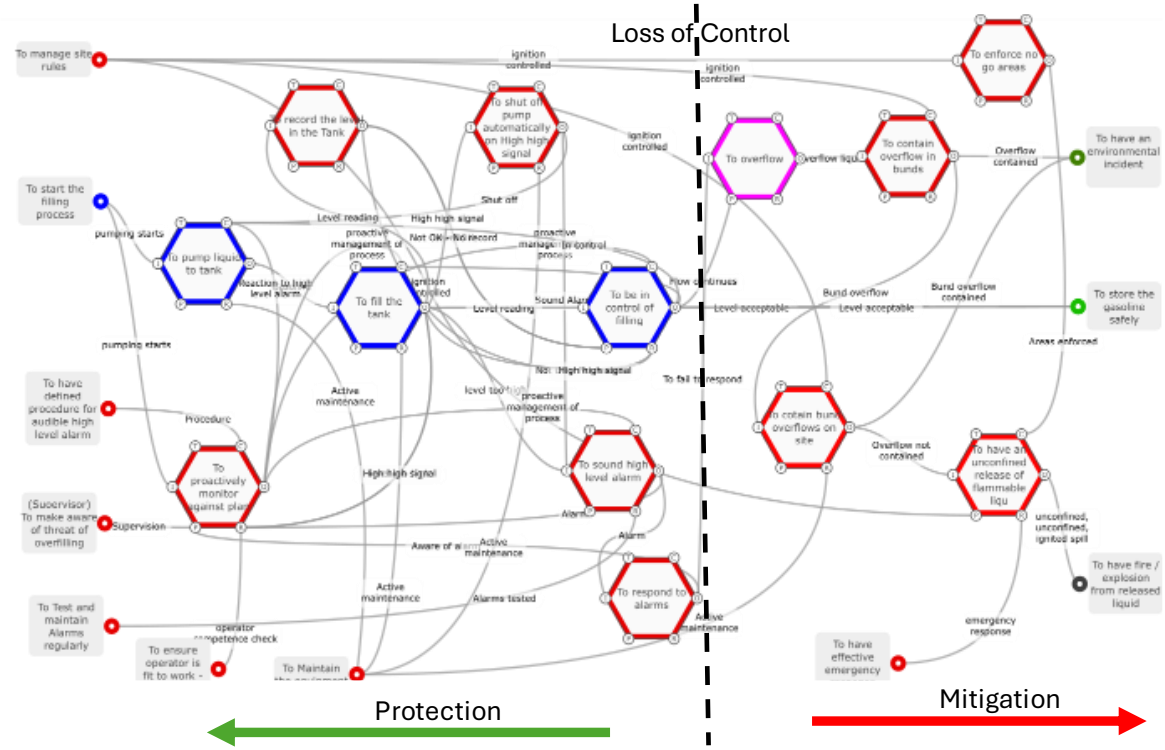

Fig. 2 – The corresponding FRAM built model for the Buncefield Barriers

With this methodology we can track the probabilities of successive functions functioning successfully. The metadata facility in the FRAM methodology allows the calculation of theses probabilities of success for each function as the process progresses. The final outcome is thus the probability of success of the process the system is designed to facilitate. Obviously then the complement of this is the probability of system failure.

So, from this FRAM built system model we can calculate the combined probability of the success or failure of the operation of these functional barriers. In the metadata algorithms we have a choice of using LOPA or Bayesian equations to calculate these results, but having the full system map, it is probably not necessary to use the Bayesian method as using the probabilities of the linking FRAM Aspects provides the interdependency evidence that the Bayes theorem uses to modify the predicted probabilities, necessary.

This has been done in the FRAM model with the results shown in Figure 3.

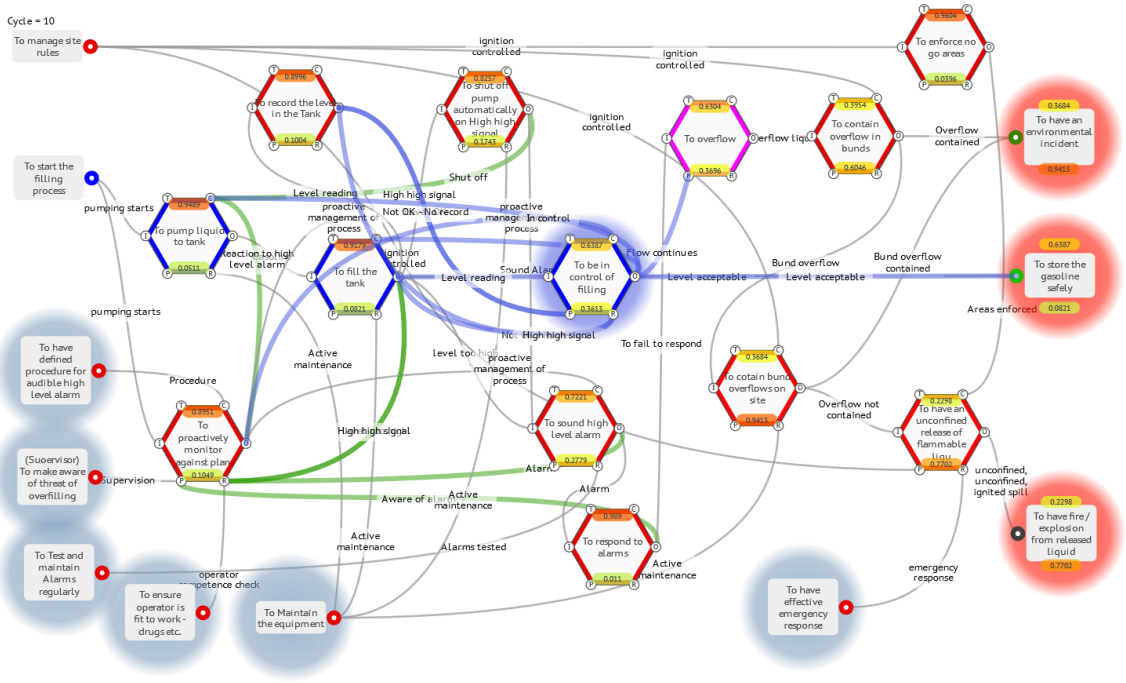

Fig. 3 – The predicted performance probabilities for the Buncefield Barriers

So, the different approaches give very different results as shown in the table below. Which set of results is more helpful and appropriate? Perhaps the lessons of Buncefield can better inform our answer?

Table 1 – Comparison of system integrity levels from various perspectives

Barrier	LOPA PFD	FRAM Pfd	Actual PFD	Note
Level recorder -	10-3	0.1	1	Frozen – must have happened before?
Hi level Switch –	10-3	0.17	1	Inoperative
Alarms –	10-2	0.27	0.5	Assumed overlooked or ignored some of the time?)
To respond to alarms -	10-1	0.01	1	Expected response to risk?
To Monitor –	10-1	0.1	1	
To have an overflow -		0.36	1	The probability of loss of control – (The BOW)
To contain –	0.1	0.9	1	Bund seals were not fireproof
Ignition control –	0.1	0.04	1	Assume only on plant situations considered

Firefighting –	0.1	0.01	0.1	Assumed well drilled
Fire / explosion if spill –	0.1	0.77	1	UVCE unexpected (A Back Swan?)
System Integrity Levels's -	10-14	0.36	1	

CONCLUSION

This paper argues that we need a more accountable and auditable means of checking that these systems are in order and will in fact add the requisite level of safety and security in operation required by companies and communities. The LOPA approach treating Barrier functions as separate, independent, linearly sequential components, can, perhaps provide a false sense of this security?

The paper argues that an alternative system thinking, systems modelling approach is perhaps more appropriate, particularly as our systems get more and more complex. This approach would allow us to capture the more subtle and real-life interactions and interdependencies between the functions, particularly those that, as happens in real life, are less than perfect and come to be ignored, or worked around. So LOPA and SIL's are designed to assure compliance with safety rules and standards for "Work as Imagined" or ideal conditions. The alternative quantitative metamodelling FRAM approach is better suited to be able to deal with real complex systems under real world "Work as Done", warts and all, conditions. It depends on the objective of the analysis, just convenience and compliance, or responsible resilience in designing out problems not just erecting barriers to prevent the inevitable.

The Buncefield disaster offers valuable lessons for managing safety in high-risk industries. One of the key takeaways is the need for redundancy in safety systems. Relying on a single level gauge or safety switch increases vulnerability to failure, especially in scenarios involving high-stakes operations like fuel storage. Regular maintenance and testing of safety-critical devices are essential to ensure their reliability. Training for operators must emphasize the importance of manual verification and proactive responses to anomalies, rather than over-relying on automation. Alarm systems must be designed to escalate effectively, ensuring that critical warnings cannot be overlooked. Secondary containment systems, such as bunds, must be robust enough to handle worst-case scenarios, with maintenance and capacity reviews conducted regularly.

The Buncefield disaster remains a stark reminder of the catastrophic potential of cascading safety failures. While technical malfunctions initiated the sequence of events, human and organizational factors allowed the situation to escalate unchecked. A holistic approach to safety, emphasizing redundancy, proactive maintenance, and a strong safety culture, is crucial in preventing similar incidents. As industries evolve, integrating advanced monitoring systems with improved training and accountability will be vital in safeguarding against disasters of this magnitude. In this way, Buncefield stands as a sombre yet invaluable case study in the importance of learning from past failures to ensure a safer future.

REFERENCES

1. Buncefield Major Incident Investigation Board. (2008). *The Buncefield Incident 11 December 2005: The Final Report of the Major Incident Investigation Board*. Health and Safety Executive. Retrieved from https://www.icheme.org/media/13707/buncefield-miib-final-report-volume-1.pdf

2. Duchene, F., Cassini, P., & Suhr, M. (2007). *Lessons learned from Toulouse and Buncefield disasters: From risk analysis failures to the identification of atypical scenarios through a better knowledge management*. Retrieved from https://www.academia.edu/5501912/Lessons_learned_from_Toulouse_and_Buncefield_disasters

3. CCPS/EI (Centre for Chemical Process Safety / Energy Institute). (2014). *Standardisation of Bow Tie Methodology and Terminology via a CCPS/EI Book*. IChemE. Retrieved from https://www.icheme.org/media/15543/poster-09.pdf

4. Slater, D. (2000), *Lessons from Flixborough*, Conference: Presentation to the 40th UK Explosion Liaison Group Meeting 25th Anniversary Meeting -19th 20th September 2000, DOI: 10.13140/RG.2.1.2764.5689

5. Exida. (n.d.). *Bow Ties Part II: Do Bow Ties have a place in Alarm Management?*. Retrieved from https://www.exida.com/Blog/bowties-part-ii-do-bow-ties-have-a-place-in-alarm-management

6. Ackoff, R. L. (1971). *Towards a System of Systems Concepts*. Management Science, 17(11), 661–671.

7. Centre for Chemical Process Safety (CCPS). (2001). *Layers of Protection Analysis: Simplified Process Risk Assessment*. Wiley-AIChE. Retrieved from https://www.aiche.org/ccps/publications/books/layers-protection-analysis-simplified-process-risk-assessment

8. Layfield, F. (1987). *Sizewell B Public Inquiry: Report by Sir Frank Layfield*. Her Majesty's Stationery Office. UK

9. Slater, D. and Hill, R. (2024), Building Nonlinear, Systemic Bow Ties, Using Functional Barriers, doi: 10.20944/preprints202406.1433.v1 (accessed 3 12 24 through https://www.researchgate.net/publication/381583186_Building_nonlinear_systemic_Bow_Ties_using_Functional_Barriers)

10. Hill, R, and Slater, D. (2024), HOW TO USE THE METADATA FACILITY IN FRAM, FRAMsynt 9/24, (Accessed 3 12 24 through https://www.researchgate.net/publication/384763526_HOW_TO_USE_THE_METADATA_FACILITY_IN_FRAM)

APPENDIX A - THE PFD PARADOX

If we multiply two probabilities of failure on demand together, we get a reduced probability of failure as an outcome. But if we multiply two probabilities of successfully operating, we get a reduced overall probability of success. How is this possible?

1. Probability of Failure on Demand (PFD)

The **Probability of Failure on Demand (PFD)** represents the likelihood that a system or component will fail when it is required to operate.

- If you have two independent systems A and B with PFDs $P_F(A)$ and $P_F(B)$, the combined system's PFD is calculated as the probability that **both systems fail** simultaneously:

$$P_{\text{F,combined}} = P_F(A) \times P_F(B)$$

Since $P_F(A)$ and $P_F(B)$ are typically small values (e.g., 0.01 or less), their product will be even smaller, meaning the combined probability of failure is reduced. This is the principle of redundancy: adding independent systems reduces the likelihood of a complete failure.

2. Probability of Success

The **Probability of Success (P(S))** is the likelihood that a system or component will operate successfully when required. If the system has a probability of failure P_F, the probability of success is:

$$P_S = 1 - P_F$$

For two independent systems A and B, the probability that **both systems succeed** simultaneously is:

$$P_{\text{S,combined}} = P_S(A) \times P_S(B)$$

Since $P_S(A)$ and $P_S(B)$ are usually less than 1 (e.g., 0.99), their product will always be smaller than either $P_S(A)$ or $P_S(B)$. This reduction reflects the reality that the chance of both systems working together is inherently lower than the chance of either system working individually.

3. Why This Seems Paradoxical

The key is understanding that probabilities are inherently multiplicative when dealing with **independent events**. Here's why the outcomes differ:

- **When multiplying probabilities of failure**, you're calculating the chance that two rare events (failures) happen simultaneously, which is rarer than either failure happening alone.
- **When multiplying probabilities of success**, you're calculating the chance that two likely events (successes) happen simultaneously, which is less likely than either success occurring individually.

4. Intuitive Example

Imagine two lightbulbs, each with a 99% chance of working correctly when switched on:

- The probability of both bulbs working together is:

$$0.99 \times 0.99 = 0.9801\ (98.01\%)$$

  This is slightly lower than 99%, as expected.

- Now consider the probability that both fail, which corresponds to the product of their failure probabilities:

$$0.01 \times 0.01 = 0.0001\ (0.01\%)$$

  This is much smaller than the failure probability of a single bulb.

5. Conclusion

The results are not contradictory but reflect the mathematical nature of probability. Multiplying probabilities of independent events reduces the likelihood of joint occurrence—whether it's two successes or two failures. For failures, the reduction is beneficial because it reflects improved reliability; for successes, it shows the added risk of requiring multiple components to work flawlessly together.

Now if we use Bayesian equation for finding the probability of success it is increased?

Yes, using Bayes' theorem can indeed increase the **probability of success** under certain conditions. This happens because Bayes' theorem incorporates **additional information or evidence** into the calculation, which can refine the estimated probability. Let's explore why and how this works.

1. Bayes' Theorem Recap

Bayes' theorem calculates the probability of an event A (e.g., success) given some observed evidence B. It is expressed as:

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

Where:

- $P(A)$: Prior probability of A (e.g., the baseline probability of success).
- $P(B|A)$: Likelihood of observing B if A is true (e.g., the reliability of the system under successful operation).
- $P(B)$: Marginal probability of B, the evidence (e.g., overall chance of observing the system behavior).

2. Why Does Probability of Success Increase?

The key to increasing the probability of success lies in the **evidence** used in B. When evidence supports success (e.g., observations indicating the system works as expected), Bayes' theorem updates the prior probability $P(A)$ to a higher posterior probability $P(A|B)$. This reflects an improved confidence in the system's success based on new, favorable evidence.

3. Example of Improved Probability of Success

Imagine a system with a prior probability of success $P(S) = 0.8$ (80%). Now, suppose we conduct a test where the system works successfully, and the test is highly reliable, meaning $P(\text{test success}|\text{system success}) = 0.95$, and the chance of the test being positive regardless of success is $P(\text{test success}) = 0.85$.

Using Bayes' theorem:

$$P(S|\text{test success}) = \frac{P(\text{test success}|S) \cdot P(S)}{P(\text{test success})}$$

Substituting the values:

$$P(S|\text{test success}) = \frac{0.95 \cdot 0.8}{0.85} = 0.894$$

The probability of success increases from $P(S) = 0.8$ to $P(S|\text{test success}) = 0.894$ (89.4%) after incorporating the test evidence.

4. What Causes the Increase?

The probability increases because the evidence P(B) (e.g., a successful test) provides new information that supports the system's success. If the test is designed to be highly sensitive and specific (reliable for identifying success), the updated probability will significantly boost confidence in success.

5. When Would It Not Increase?

The probability of success might not increase if:

- The evidence P(B) is neutral or does not strongly correlate with success.
- The evidence suggests failure instead (e.g., an unreliable or failed test).
- The prior probability is already very low, and the evidence is weak, so the posterior probability remains small.

6. Implications for System Design

In practical terms, Bayesian updates are extremely useful for dynamic systems where:

- You can incorporate operational feedback, test results, or sensor data.
- Decisions rely on continuously updated confidence in the system's success.

For example, in reliability engineering, Bayesian methods are used to adjust predictions of system performance based on observed operational data. This dynamic adjustment helps engineers refine maintenance schedules, improve redundancy planning, and better anticipate failures.

7. Conclusion

Bayesian methods often increase the probability of success because they integrate evidence favouring success into the calculation, refining the confidence in the outcome. This makes Bayesian analysis a powerful tool for improving predictions and managing uncertainty in complex systems.