

Redundancy Buys Less Than It Promises: Common-Cause Saturation and the Value of Proof Testing in Safety-Instrumented Architectures

Dr. Jayant Damodar Supe

Associate Professor, Department of Civil Engineering, RCET, Bhilai, Chhattisgarh, India.

Email: jayantcd@gmail.com

ABSTRACT

Redundant voting architectures are the standard means of achieving demanding safety integrity targets. Their calculated benefit depends on the assumption that channels fail independently, which they do not: shared design, shared environment, shared calibration and shared maintenance produce correlated failures represented in practice by a common-cause factor. How far that factor erodes the benefit of redundancy, and where the remaining design leverage lies, is quantified here. This study determines how the probability of failure on demand of 1oo1, 1oo2, 2oo2, 2oo3 and 1oo3 architectures varies with the common-cause factor and the proof-test interval, identifies the point at which additional redundancy ceases to pay, and compares redundancy against single-channel improvement and more frequent testing on a common basis. Simplified low-demand probability of failure on demand relations were evaluated for a dangerous undetected channel failure rate of 0.02 per year, proof-test intervals from 0.25 to 5 years, and common-cause factors from 0 to 20 %. The share of the resulting probability attributable to common cause was computed for each architecture, together with the equivalent single-channel improvement needed to match a redundant architecture and the spurious trip rate of each configuration. With independent channels, 1oo3 reduced the probability of failure on demand by a factor of 5,000 relative to a single channel. At a common-cause factor of 1 % that factor fell to 98, at 5 % to 19.9, and at 10 % to 10.0. In the same conditions 1oo2 delivered 16.1 and 1oo3 delivered 19.9, so the third channel bought a 24 % improvement for 50 % more hardware. At a 5 % common-cause factor, 98.4 % of the residual failure probability of a 1oo3 architecture was common cause. Shortening the proof-test interval from 5 years to 0.25 years reduced the 1oo2 probability by a factor of 42, exceeding the range spanned by any architectural choice at that common-cause factor. Matching a 1oo2 architecture with a single channel required a 16.1-fold improvement in channel failure rate. Beyond duplication, additional channels deliver diminishing returns that are governed almost entirely by the common-cause factor rather than by the voting logic. Where the common-cause factor exceeds a few per cent, effort spent on diversity, on proof-test frequency or on channel quality dominates effort spent on further redundancy.

Keywords: functional safety; common-cause failure; redundancy; voting architectures; probability of failure on demand; systems engineering.

1. INTRODUCTION

When a safety function must be more reliable than a single instrument can be, the standard answer is to use more than one [1,2,10]. Two transmitters voting one-out-of-two, three

voting two-out-of-three, or triple modular redundancy in a controller are the familiar architectures, and the arithmetic supporting them is compelling: if a channel fails with probability p and channels fail independently, two channels both fail with probability p^2 [6,8].

Independence is the assumption doing all the work, and it is never quite true. Redundant channels are usually the same model from the same supplier, built to the same design, installed by the same team on the same day, exposed to the same process fluid and the same ambient conditions, calibrated with the same procedure and the same reference, and maintained by the same technician working from the same instructions. A systematic fault in any of those is a fault in all channels at once.

Functional safety practice handles this with a common-cause factor — a fraction of channel failures treated as affecting all channels simultaneously [3-5]. The device is crude and it is the best available, and its consequence is not always appreciated: because the common-cause term is proportional to the single-channel failure probability rather than to its square or cube, it eventually dominates, and once it does, further redundancy achieves almost nothing.

This study quantifies where that happens. Four questions are addressed: how the probability of failure on demand of the standard architectures varies with the common-cause factor; what fraction of the residual risk is common cause in each; how redundancy compares against the two alternatives available to a designer, namely a better channel and a shorter proof-test interval; and what the safety gains cost in spurious trips.

2. MODEL

2.1 Probability of Failure on Demand

For a low-demand safety function proof-tested at interval T with a dangerous undetected channel failure rate λ , the average probability of failure on demand for a single channel is $\lambda T/2$ [7]. Redundant architectures are evaluated with the standard simplified relations [1,6,8,9], splitting each channel's failure rate into an independent part $(1 - \beta)\lambda$ and a common-cause part $\beta\lambda$:

$$PFD_{1oo2} = \frac{[(1 - \beta)\lambda T]^2}{3} + \frac{\beta\lambda T}{2} \quad (1)$$

$$PFD_{2oo3} = [(1 - \beta)\lambda T]^2 + \frac{\beta\lambda T}{2} \quad (2)$$

$$PFD_{1oo3} = \frac{[(1 - \beta)\lambda T]^3}{4} + \frac{\beta\lambda T}{2} \quad (3)$$

The structure of these expressions is the whole argument. The first term falls as a power of the number of channels; the second does not fall at all. Whatever the voting logic, the common-cause term is the same $\beta\lambda T/2$, so it places a floor under the achievable probability that no amount of redundancy penetrates. The model parameters are listed in Table 1.

2.2 Quantities Computed

Four quantities were derived. The probability of failure on demand for each architecture across the common-cause and test-interval ranges; the improvement factor relative to a single channel; the share of the resulting probability contributed by the common-cause term; and the single-channel failure rate that would be needed to match a redundant architecture without

redundancy. Spurious trip rates were computed from the safe failure rate and the voting logic, treating a 1ooN architecture as tripping on any channel and a 2oo3 as requiring two [6,14].

Table 1: Model parameters.

Parameter	Value	Basis
Dangerous undetected channel failure rate, λ	0.02 per year	Representative of a well-characterised field instrument
Proof-test interval, T	1 year baseline; 0.25 to 5 years explored	Typical range in process industry practice
Common-cause factor, β	0 to 20 %	Spans idealised independence to poor diversity
Safe failure rate per channel	0.1 per year	Used only for the spurious trip comparison
Architectures	1oo1, 1oo2, 2oo2, 2oo3, 1oo3	Standard voting configurations
Repair time	Neglected	Low-demand mode; undetected failures dominate

Note: Values are representative assumptions rather than data for any device. The relations used are the conventional simplified forms; they neglect diagnostic coverage, repair during the test interval and non-perfect proof testing, as discussed in Section 4.3.

3. RESULTS

3.1 The Effect of Common Cause

Table 2 and Figure 1 give the probability of failure on demand of each architecture across the range of common-cause factors.

The 1oo3 column is the striking one. With perfectly independent channels it achieves 2.00×10^{-6} , comfortably within the most demanding integrity band. A common-cause factor of one per cent — a very good installation — raises that to 1.02×10^{-4} , a fiftyfold degradation, and at five per cent it reaches 5.02×10^{-4} .

Table 2: Probability of failure on demand by architecture and common-cause factor.

β	1oo1	1oo2	2oo2	2oo3	1oo3
0 %	1.00×10^{-2}	1.33×10^{-4}	2.00×10^{-2}	4.00×10^{-4}	2.00×10^{-6}
1 %	1.00×10^{-2}	2.31×10^{-4}	2.00×10^{-2}	4.92×10^{-4}	1.02×10^{-4}
2 %	1.00×10^{-2}	3.28×10^{-4}	2.00×10^{-2}	5.84×10^{-4}	2.02×10^{-4}
5 %	1.00×10^{-2}	6.20×10^{-4}	2.00×10^{-2}	8.61×10^{-4}	5.02×10^{-4}
10 %	1.00×10^{-2}	1.11×10^{-3}	2.00×10^{-2}	1.32×10^{-3}	1.00×10^{-3}
20 %	1.00×10^{-2}	2.09×10^{-3}	2.00×10^{-2}	2.26×10^{-3}	2.00×10^{-3}

By that point 1oo3 and 1oo2 differ by less than a factor of 1.3, and both have converged toward the common-cause floor. The 2oo2 column is included as a reminder that not all redundancy improves safety: requiring both channels to demand a trip doubles the probability of failure on demand relative to a single channel, because either channel failing dangerously defeats the function.

3.2 How Quickly the Benefit Saturates

Table 3 and Figure 2 give the improvement factor of each architecture relative to a single channel.




Fig.1: Probability of failure on demand against the common-cause factor, with integrity level boundaries marked.

Table 3: Improvement factor relative to a single channel.

β	1oo2	2oo3	1oo3	1oo3 gain over 1oo2
0 %	75.0	25.0	5,000.0	66.7×
1 %	43.4	20.3	98.1	2.26×
2 %	30.5	17.1	49.5	1.62×
5 %	16.1	11.6	19.9	1.24×
10 %	9.0	7.6	10.0	1.11×




Fig.2: Improvement factor over a single channel against the common-cause factor.

The final column is the design question in one number. With independent channels, adding a third channel to a duplex arrangement improves the probability of failure on demand by a factor of 66.7. At a common-cause factor of one per cent that becomes 2.26; at five per cent, 1.24; at ten per cent, 1.11.

In other words, for any installation with a common-cause factor above a few per cent, the third channel delivers a ten to twenty-five per cent improvement in exchange for fifty per cent more hardware, fifty per cent more installation, fifty per cent more calibration and fifty per cent more maintenance — all of which are themselves sources of common-cause failure [12]. It is entirely possible for a third channel to raise the effective common-cause factor by more than it lowers the independent term, in which case the additional redundancy makes the function worse.

3.3 What the Residual Risk Consists Of

Table 4 and Figure 3 give the share of the probability of failure on demand that arises from common cause.

Table 4: Share of the probability of failure on demand arising from common cause.

β	1oo2	2oo3	1oo3
1 %	43.4 %	20.3 %	98.1 %
2 %	61.0 %	34.2 %	99.1 %
5 %	80.6 %	58.1 %	99.7 %
10 %	90.3 %	75.5 %	99.9 %




Fig.3: Proportion of the residual failure probability attributable to common cause.

For a 1oo3 architecture, between 98.1 % and 99.9 % of the residual probability of failure on demand is common cause across the whole range examined. Triple redundancy does not have an independent-failure problem; it has a common-cause problem and nothing else.

This reframes where verification effort should go. A design review that scrutinises channel failure rates, voting logic and diagnostic coverage for a 1oo3 architecture is examining the terms that contribute one part in a hundred. The quantity that determines the result is the common-cause factor, which is estimated rather than measured, is usually assigned from a checklist [1,12], and is the least defensible number in the calculation.

The 2oo3 column shows a lower common-cause share than 1oo2 at every value, which at first looks like an advantage. It is not: the share is lower because the independent term is larger, so 2oo3 has a higher total probability than 1oo2 throughout. The architecture is chosen for its spurious trip behaviour rather than its safety performance, as Section 3.5 shows.

3.4 Redundancy against the Alternatives

Table 5 and Figure 4 give the probability of failure on demand against the proof-test interval at a 5 % common-cause factor.

Table 5: Probability of failure on demand against proof-test interval at $\beta = 5\%$.

Test interval	1oo1	1oo2	2oo3	1oo3
3 months	2.50×10^{-3}	1.33×10^{-4}	1.48×10^{-4}	1.25×10^{-4}
6 months	5.00×10^{-3}	2.80×10^{-4}	3.40×10^{-4}	2.50×10^{-4}
1 year	1.00×10^{-2}	6.20×10^{-4}	8.61×10^{-4}	5.02×10^{-4}
2 years	2.00×10^{-2}	1.48×10^{-3}	2.44×10^{-3}	1.01×10^{-3}
5 years	5.00×10^{-2}	5.51×10^{-3}	1.15×10^{-2}	2.71×10^{-3}




Fig.4: Probability of failure on demand against proof-test interval at a 5 % common-cause factor.

Shortening the proof-test interval from five years to three months improved the 1oo2 probability by a factor of 41.5 [17]. Across the same conditions, moving from 1oo2 to 1oo3 improved it by a factor of 1.24. The maintenance schedule is worth more than the architecture by a factor of thirty, and it is a cheaper decision to reverse.

The comparison with channel quality points the same way. Matching a 1oo2 architecture at a 5 % common-cause factor with a single channel requires a dangerous undetected failure rate of 1.24×10^{-3} per year against the baseline 2×10^{-2} — a 16.1-fold improvement. That is a demanding target for a field device and it is not absurd; it is roughly the difference between a general-purpose transmitter and a certified, well-diagnosed one [13], achieved with no additional wiring, no additional calibration and no additional common-cause exposure.

Together these results suggest an ordering of design effort that inverts common practice: establish the test interval and the channel quality first, add duplication where the target still is not met, and treat triplication as a decision about availability rather than about safety.

3.5 The Availability Trade

Table 6 and Figure 5 set the spurious trip rate of each architecture against its safety performance.

Table 6: Spurious trip rate against safety performance.

Architecture	PFD at $\beta = 5\%$	Spurious trips per year	Comment
1oo1	1.00×10^{-2}	0.10	Baseline on both measures
1oo2	6.20×10^{-4}	0.20	Best safety of the duplex options, worst availability
2oo3	8.61×10^{-4}	0.03	Slightly worse safety, three times better availability than 1oo1
1oo3	5.02×10^{-4}	0.30	Marginal safety gain, worst availability




Fig.5: Spurious trip rate against probability of failure on demand for four architectures.

The 2oo3 architecture is the interesting case. On safety alone it is worse than 1oo2 by a factor of 1.4, which would not on its own justify the extra channel. On availability it is better than 1oo1 by a factor of 3.3 and better than 1oo2 by a factor of 6.7, because two channels must fail spuriously before the plant trips.

This is why 2oo3 predominates in practice and it is worth being explicit about the reason. It is not chosen because it is the safest arrangement, which it is not, but because it is the arrangement that delivers most of the safety benefit of duplication without the process upsets that 1oo2 causes. Framing 2oo3 as a safety choice rather than an availability choice obscures the trade actually being made, and leads to its selection in applications where a spurious trip is harmless and 1oo2 would be better.

4. DISCUSSION

4.1 Principal Findings

Four results follow. The benefit of redundancy saturates rapidly in the common-cause factor: 1oo3 improves on a single channel by 5,000-fold with independent channels, 98-fold at $\beta = 1\%$ and 19.9-fold at $\beta = 5\%$. Beyond duplication the returns are small, with the third channel worth 24% at $\beta = 5\%$ for 50% more hardware. The residual probability of a 1oo3 architecture is 98–99.9% common cause, so the parameters usually scrutinised contribute almost nothing. And the proof-test interval spans a factor of 41.5 against the architecture's 1.24, while a 16.1-fold better single channel matches 1oo2 outright.

4.2 Implications for Practice

Table 7 summarises each finding, its implication and the corresponding recommendation.

Table 7: Implications.

Finding	Implication	Recommendation
Benefit saturates in β	Independence assumptions dominate the result	Report PFD against a range of β , never a single value
Third channel worth 24 % at $\beta = 5$ %	Triplication is poor value for safety	Justify a third channel on availability, not on PFD
1oo3 residual is 98–99.9 % common cause	Channel data barely affect the answer	Direct verification effort at the β estimate and its basis
Test interval worth 41.5×	Maintenance beats architecture	Fix the test interval before choosing the voting logic
16.1× better channel matches 1oo2	Quality substitutes for redundancy	Price a certified channel against a second channel and its lifetime cost
2oo2 is worse than 1oo1	Not all redundancy is protective	Check that the voting logic serves the failure mode of concern
2oo3 worse than 1oo2 on PFD, better on trips	It is an availability architecture	Select it where spurious trips are costly, not where safety is tightest

The third row carries the most practical weight. In a typical verification, channel failure rates are taken from manufacturer data with several significant figures, diagnostic coverage is argued at length, and the common-cause factor is assigned from a scoring checklist in an afternoon. The results here show that ordering to be exactly inverted relative to the influence each has on the answer, and the imbalance is largest for the architectures used in the most critical applications.

The second row has a corollary that deserves stating. Because each additional channel brings its own installation, calibration and maintenance activity, and because those activities are themselves common-cause mechanisms, the common-cause factor is not independent of the architecture. A model that holds β fixed while adding channels, as this one does, is generous to redundancy; if β rises with channel count, the case for triplication weakens further and can invert.

4.3 Limitations

The simplified probability relations used here neglect several effects present in a full treatment [11,15,16]. Diagnostic coverage and repair of detected failures within the test interval are omitted, which raises all reported probabilities relative to a system with good diagnostics but affects the architectures similarly and does not change their ordering. Imperfect proof testing — the fraction of failures a test does not reveal — is not modelled and would place an additional floor under the benefit of more frequent testing, weakening the strongest recommendation made here.

The beta-factor model of common cause is itself a coarse approximation. It assumes that a common-cause event fails all channels simultaneously and completely, whereas real dependent failures are partial and staggered, and more elaborate models such as the multiple Greek letter or alpha-factor approaches distinguish the failure of two channels from three [4,5,14]. For 1oo3 in particular the beta model is known to be conservative in some conditions and optimistic in others, so the extreme common-cause shares reported in Table 4 should be read as indicating dominance rather than as precise values.

As noted above, β is held constant across architectures, which favours redundancy. Sensor and final-element subsystems are not distinguished, and a real safety function combines several subsystems whose probabilities add, so the achievable integrity of a complete function is worse than any figure here. The spurious trip comparison uses a single safe failure rate and ignores repair, and treats 2oo2 as effectively never tripping spuriously, which overstates its availability.

Finally, all parameter values are representative assumptions rather than data for a specific device or installation. Absolute probabilities should not be quoted for any real system, and the results are offered as a structural account of how the terms interact rather than as a verification of any architecture.

5. CONCLUSION

The case for redundancy rests on independence, and the common-cause factor measures how far independence fails. On the model examined, a 1oo3 architecture improves on a single channel by a factor of 5,000 when channels are independent, 98 when one per cent of failures are common cause, and 19.9 when five per cent are — by which point 98.4 % of its residual failure probability is common cause and the third channel is worth 24 % over a duplex arrangement.

Two conclusions follow. Verification effort is misallocated when it concentrates on channel failure rates and voting logic in architectures whose residual risk is almost entirely dependent failure; the common-cause estimate is the number that determines the answer and is usually the least examined. And the alternatives to redundancy are stronger than they are usually credited: quartering the proof-test interval was worth a factor of 41.5 where a third channel was worth 1.24, and a sixteenfold better single channel matched a duplex arrangement outright, without adding any of the shared installation and maintenance that generates common-cause failure in the first place.

DECLARATIONS

Ethical approval. Not applicable. This is a computational study; no human or animal subjects were involved.

Funding. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Conflicts of interest. The author declares no conflicts of interest, financial or non-financial, in relation to the subject matter of this manuscript.

Scope disclaimer. This paper is a structural analysis of reliability relations and is not a functional safety verification. It does not constitute compliance advice, and the simplified relations used are not a substitute for the requirements of the applicable standard. No device, vendor or installation is evaluated.

Parameter disclaimer. All values in Table 1 are representative assumptions. Absolute probabilities apply only to the case modelled and must not be cited for any real safety function.

Data and code availability. This is a computational study; no experimental or field data were used. All results follow from Equations (1) to (3) using the parameters listed in Table 1, and every reported value can be reproduced from them. The code is available from the corresponding author upon reasonable request.

Author's contribution. Dr. Jayant Damodar Supe is the sole author of this work and was responsible for the conception and design of the study, development of the reliability model, analysis and interpretation of the results, preparation of the figures and tables, and drafting,

revision and final approval of the manuscript. The author is accountable for all aspects of the work.

REFERENCES

1. International Electrotechnical Commission. IEC 61508: Functional safety of electrical/electronic/programmable electronic safety-related systems. Geneva: IEC; 2010.
2. International Electrotechnical Commission. IEC 61511: Functional safety — safety instrumented systems for the process industry sector. Geneva: IEC; 2016.
3. Fleming KN. A reliability model for common mode failure in redundant safety systems. General Atomic Report GA-A13284; 1975.
4. Mosleh A, Fleming KN, Parry GW, et al. Procedures for treating common cause failures in safety and reliability studies. NUREG/CR-4780. Washington, DC: US Nuclear Regulatory Commission; 1988.
5. Hokstad P, Rausand M. Common cause failure modeling: status and trends. In: Handbook of Performability Engineering. London: Springer; 2008:621–640.
6. Rausand M. Reliability of Safety-Critical Systems: Theory and Applications. Hoboken: Wiley; 2014.
7. Rausand M, Høyland A. System Reliability Theory: Models, Statistical Methods, and Applications. 2nd ed. Hoboken: Wiley; 2004.
8. Goble WM, Cheddie H. Safety Instrumented Systems Verification: Practical Probabilistic Calculations. Research Triangle Park: ISA; 2005.
9. Smith DJ. Reliability, Maintainability and Risk. 9th ed. Oxford: Butterworth-Heinemann; 2017.
10. Summers AE. Introduction to layers of protection analysis. J Hazard Mater. 2003;104(1–3):163–168.
11. Jin H, Lundteigen MA, Rausand M. Reliability performance of safety instrumented systems: a common approach for both low- and high-demand mode of operation. Reliab Eng Syst Saf. 2011;96(3):365–373.
12. Lundteigen MA, Rausand M. Common cause failures in safety instrumented systems on oil and gas installations: implementing defense measures through function testing. J Loss Prev Process Ind. 2007;20(3):218–229.
13. Brissaud F, Barros A, Bérenguer C, Charpentier D. Reliability analysis for new technology-based transmitters. Reliab Eng Syst Saf. 2011;96(2):299–313.
14. Hauge S, Håbrekke S, Lundteigen MA. Reliability prediction method for safety instrumented systems — PDS method handbook. SINTEF Report A24442; 2013.
15. Zhang T, Long W, Sato Y. Availability of systems with self-diagnostic components — applying Markov model to IEC 61508-6. Reliab Eng Syst Saf. 2003;80(2):133–141.
16. Bukowski JV. Incorporating process demand into models for assessment of safety system performance. Proc Annu Reliab Maintainab Symp. 2006:577–581.
17. Torres-Echeverría AC, Martíorell S, Thompson HA. Modelling and optimization of proof testing policies for safety instrumented systems. Reliab Eng Syst Saf. 2009;94(4):838–854.